Bank of Canada Banque du Canada

Working Paper 2003-26 / Document de travail 2003-26

Measuring Interest Rate Expectations in Canada

by

Grahame Johnson

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2003-26

September 2003

Measuring Interest Rate Expectations in Canada

by

Grahame Johnson

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
gjohnson@bankofcanada.ca

The views expressed in this paper are those of the author.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements iv
Abstract/Résumé v

1. Introduction 1

2. The Expectations Hypothesis 3
   2.1 Definition 3
   2.2 Empirical evidence 4

3. Definition and Selection of Instruments 7
   3.1 Definition of instruments 7
   3.2 Selection of instruments 11

4. A Comparison of Pre- and Post-FAD Periods 14
   4.1 Statistical summary of the periods 14
   4.2 Testing the expectations hypothesis 20
   4.3 Financial rankings and term premium estimates 27

5. Expectations Over a Longer Horizon 30
   5.1 Selection of instruments 30
   5.2 Testing bankers' acceptance futures 31
   5.3 Generation of a spot BA yield curve 35

6. Implied Forward Rates 36

7. Conclusions 43

Bibliography 44

Acknowledgements

I would like to thank Ron Morrow for guiding this work along from the beginning. As well, thanks to Scott Hendry, Chris D'Souza, Jing Yang, and Dave Bolder for many useful comments and much-needed guidance. Any errors, omissions, or inconsistencies are my own.

Abstract

Financial market expectations regarding future policy actions by the Bank of Canada are an important input into the Bank's decision-making process, and they can be measured using a variety of sources. The author develops a simple expectations-based model to focus on measuring interest rate expectations that are implied by the current level of money market yields. The explanatory power of this model increases markedly in the period following the implementation of the Bank's regime of fixed announcement dates in November 2000, and it appears to accurately describe the behaviour of short-term yields. Term premiums are estimated for the various instruments examined, and observed market yields are adjusted by those amounts. Once the market yields are adjusted, they can be used to calculate implied forward rates for a series of dates in the future. These forward rates can be interpreted as representing the market's expectations for the future level of overnight rates at a specific date.

JEL classification: G1
Bank classification: Financial markets; Interest rates

Résumé

Les attentes des marchés financiers concernant les décisions futures de la Banque du Canada en matière de politique monétaire occupent une place importante dans le processus décisionnel de l'institution, et on peut les mesurer à l'aide de diverses sources. L'auteur élabore un modèle simple pour mesurer spécifiquement les attentes de taux d'intérêt implicites dans les rendements courants du marché monétaire. Le modèle voit son pouvoir explicatif s'accroître sensiblement dans la période qui a suivi la mise en œuvre, en novembre 2000, du système de dates d'annonce préétablies et il semble décrire avec précision le comportement des rendements à court terme. L'auteur évalue les primes à terme pour les divers instruments examinés, puis corrige en conséquence les rendements observés. Une fois corrigés, les rendements du marché peuvent servir à calculer les taux à terme implicites pour une série de dates à venir. On peut considérer ces taux à terme comme représentatifs des attentes des marchés concernant le niveau auquel le taux du financement à un jour se situera à une date précise.

Classification JEL : G1
Classification de la Banque : Marchés financiers; Taux d'intérêt

1. Introduction

The measurement of financial market expectations regarding future changes in the overnight rate is an important input into the Bank of Canada's decision-making process for setting the target for the overnight rate. The Bank of Canada adjusts this target in an attempt to influence the inflation rate. The linkage between the overnight rate and the inflation rate consists of three key steps: the first is between the overnight rate and other financial variables (longer-term interest rates and the exchange rate), the second is from these financial variables to aggregate demand, and the third is from aggregate demand to the output gap. The financial markets are the mechanism through which the first step is realized and through which changes in the overnight rate are transmitted into the other financial variables. It is important, therefore, to be able to properly assess the impact that contemplated policy decisions may have on market-determined interest rates, because these interest rates feed into the real economy. To this end, policy-makers need to be aware of which decisions would constitute a surprise and which are well anticipated. An accurate measure of the market's opinion on the direction and magnitude of future rate changes can therefore help policy-makers assess the full potential impact of any contemplated rate decisions and give them a frame of reference that may influence how policy decisions are communicated.

Interest rate expectations embedded in securities prices also provide valuable information about how market participants view the economy. The short-term maturity segment of Canada's interest rate market contains a large number of liquid instruments that allow participants to structure positions based on their views of the expected future path of short rates. The observed market yields of these instruments represent a sort of consensus estimate for this future path.[1] This consensus economic forecast can be compared and contrasted with the Bank's internal economic forecasts.

An accurate measure of the market's expectations for future interest rate moves, therefore, has two significant benefits. First, from a tactical perspective, knowing what interest rate path is currently discounted in market prices will help policy-makers anticipate the near-term effect of a specific interest rate decision on other, market-determined rates. Second, the market's

1. This consensus estimate is, however, dollar weighted. Participants who place the largest amount of money at risk have the largest influence on prices. These participants may be extremely confident of their view, have very high levels of risk tolerance, or have a very large capital base. None of these factors suggests, though, that their view is any more informed or accurate than that of others with less risk capital. This does mean, however, that the perfect competition assumption common in many economic models (many small participants, none of which have significant pricing power) does not hold in this case.

expectations of future interest rate moves can serve as a type of consensus economic forecast. This forecast can then be compared with those of the Bank and of other, external forecasters.

Financial market expectations regarding future changes in the overnight rate can be measured using a variety of sources, including expectations implicit in the yields of money market instruments, surveys of private sector forecasters, published reports from investment dealers, and regular interaction with market participants. This paper describes one method by which the Bank quantifies the market's expectations using observed yields on money market assets. The focus will be on measuring expectations over a relatively short time horizon, specifically twelve months and less.

The model described in this paper is based upon the expectations hypothesis (EH), which implies that longer-term interest rates are rational estimators of future short-term interest rates. Section 2 describes the EH in detail, explaining the theory and providing a brief literature review. Section 3 reviews the various Canadian money market instruments that could be potential inputs into the models and selects those that are most suitable. Sections 4 and 5 test the EH using the selected money market assets over two distinct subperiods, and section 6 demonstrates the mechanics of the actual model.

In November 2000, the Bank implemented a policy of fixed announcement dates (FADs), changing the way monetary policy was implemented. Prior to that date, the Bank could change the overnight rate on any date. This meant that the market might be confident of the direction of interest rate changes, but very uncertain as to the timing. This uncertainty could at times make it difficult to price short-term assets and may have reduced their predictive effectiveness. With the implementation of the FAD policy, the Bank committed itself to consider changes to the overnight rate on a series of eight pre-announced dates each year. Changes between FADs, while still possible, would be made only under exceptional circumstances.[2] The goal of this change was to decrease uncertainty regarding the timing of the Bank's policy changes. This change was also expected to improve the focus on domestic circumstances, resulting in increased efficiency in the pricing of short-term assets. By testing the EH in both the pre- and post-FAD periods, it is possible to determine whether this has occurred.

2. In the press release that announced the implementation of the FADs, it was stated that the Bank would retain the option of taking action between fixed dates, although it would exercise this option only in the event of extraordinary circumstances. To date, only one change has been made between FADs: on 17 September 2001, the Bank lowered the overnight rate by 50 basis points following the 11 September 2001 terrorist attacks.

2. The Expectations Hypothesis

2.1 Definition

The EH is the best-known and most intuitive theory of the term structure of interest rates. It maintains that forward rates (or spot rates, for securities with a longer term-to-maturity) are rational estimates of future realized short rates plus a constant risk premium. This is equivalent to stating that a longer-term (single period) interest rate should be equal to the geometric average of expected future short-term rates plus a risk premium:

$$Y_t(r) = E\left[\left[\prod_{i=t}^{r}(1+R_{t+i})\right]^{\frac{1}{r}}\middle|\Omega_t\right] + \alpha_r, \qquad (1)$$

where R_t is the one-period rate at time t, Ω_t is the information set at time t, $Y_t(r)$ is the r-period term rate at time t, and α_r represents a constant risk premium that can be distinct across the different maturities, r.

There are two versions of the EH. The first, the pure EH, sets α_r equal to zero and maintains that investor expectations of future short-term interest rates are the sole determinant of long-term rates. It assumes that market participants (in aggregate) are risk-neutral and that expected returns across assets of different maturities are equal for any given investment horizon. The pure EH also implies that forward rates represent unbiased estimators of future short rates, and that any excess returns from forward rates are random errors and unforecastable.

The second version, the general EH, weakens the constraint on the pure EH slightly, allowing α_r to take on non-zero values. The shape of the yield curve in the general EH is influenced by three factors: the market's expectations for future short-term interest rates, a bond risk premium, and a convexity bias. In this version, α_r amalgamates both the bond risk premium and the convexity bias.

The bond risk premium represents a yield premium (or excess expected return) that an investor requires to hold any instrument with a longer maturity or greater credit risk than the one-period risk-free asset. This risk premium does not have to be positive and can vary across sections of the yield curve. The liquidity-preference hypothesis holds that investors are risk-averse and dislike short-term fluctuations in asset prices. As a result, longer-term bonds require higher yields to compensate for their increased price volatility. The preferred-habitat hypothesis maintains that the

risk premium is a function of supply and demand imbalances within specific maturity sectors, and therefore may actually decrease with duration. Investors who typically hold long-duration liabilities (such as life insurance companies and pension funds) may perceive longer-duration assets as less risky. If the institution's goal is to immunize its liabilities, it should be willing to pay a premium for long-duration assets that accomplish this goal. In general, the term "risk premium" is used to encompass all of these factors.

Convexity describes the curvature of the yield-to-maturity curve. All non-callable bonds have positive convexity, which means that their prices increase more for a given decline in yields than they fall for an equivalent rise in yields. Convexity is a desirable characteristic, since, all else being equal, it leads to higher returns. Convexity can be considered a type of option, because it has a larger value in markets where volatility is high. Investors are willing to pay for this, so bonds with a large amount of convexity tend to trade at lower yields than bonds of similar duration that are less convex. Long-maturity zero-coupon bonds (often having durations of over 25 years) have very high levels of convexity. The result is the tendency for the very long end of the yield curve to be flat (or even inverted). Convexity is a much more significant factor with long-maturity bonds; its impact on assets that have less than one year to maturity is negligible.

A further variation of the EH maintains that the risk premium, α_r, varies across time, or, stated another way, $\alpha_{r,t} \neq \alpha_{r,t+1}$. The evidence for a time-varying value of α_r is particularly strong for longer-maturity assets, and a number of empirical studies that reject the more strict definitions of the EH allow that it could hold if the risk premiums were time-varying. Research in this area has focused on identifying the economic variables that influence the size of the risk premiums and attempting to model their behaviour going forward.

2.2 Empirical evidence

The EH has been the subject of an enormous amount of empirical work. The results have generally been consistent in rejecting both the pure EH (rejection of the hypothesis that the risk premium is zero) and the general EH (rejection of the hypothesis that the risk premium is a constant). Schiller (1990) provides a general literature review of ten of the earlier studies and finds that all reject the EH. Campbell (1995) also rejects the hypothesis and Cox, Ingersoll, and Ross (1981) argue that it contains mutually contradictory propositions and is incompatible with any continuous-time, rational-expectations equilibrium model of the yield curve. These studies do, however, consistently maintain that the EH could hold if the risk premiums were time-varying.[3]

3. Studies that suggest the EH would hold with time-varying risk premiums include Schiller, Campbell, and Schoenholtz (1983), Fama and Bliss (1987), Froot (1989), and Schiller (1990).

Empirical tests of the EH using Canadian data are rare. Hejazi, Lai, and Yang (2000) reject the hypothesis using Canadian treasury bill data based on the existence of time-varying risk premiums. They then examine the determinants of these premiums, and find that their size is correlated with the size of the yield-curve spreads between different maturities, and sensitive to the conditional variances of U.S. macroeconomic variables. Paquette and Stréliski (1998) examine the EH in Canada using forward rate agreements and also find evidence of time-varying risk premiums.

Some recent studies, however, have emerged in defence of the EH, particularly at the short end of the yield curve. Longstaff (2000a) provides a theoretical argument in which he maintains that if fixed-income markets are incomplete, then the EH cannot be ruled out on theoretical grounds, as Cox, Ingersoll, and Ross (1981) suggest. All traditional forms of the EH can be consistent with the no-arbitrage condition. Longstaff (2000b) also finds some support for the EH in the very short end of the yield curve. Specifically, he finds that, using overnight, weekly, and monthly repo rates, the term rates are unbiased estimators of the average overnight rate realized over the period. The risk premiums in the weekly and monthly rates are very small and not significantly different from zero. Durre, Snorre, and Pilegaard (2003) test the EH on daily rates in the euro area using forward and spot rates in a cointegrated vector autoregression (VAR) model. They find evidence that supports the EH for maturities out to nine months. Cole and Reichenstein (1994) test whether eurodollar futures provide an unbiased estimate of the U.S.-dollar London Inter-Bank Offered Rate (LIBOR) at their expiration.[4] They find that the front eurodollar contract provides an unbiased forecast of LIBOR at expiry, while more distant contracts contain a risk premium that increases with the time-to-maturity of the contract.

The Bank of Canada has conducted several studies that attempt to estimate time-varying risk premiums for money market instruments. Gravelle and Morley (1998) and Gravelle, Muller, and Stréliski (1998) use both a vector-error-correction model (VECM) and a Kalman filter to estimate a time-varying parameter model of excess forward returns. Both studies find evidence of a time-varying component in the risk premiums, but conclude that, for short horizons, the EH is a reasonable characterization of the behaviour of the short end of the yield curve. During periods of relative interest rate stability and stable inflation expectations, the more complicated econometric models give results almost identical to an expectations-based measure. It is only during "crisis"

4. LIBOR is the most widely used benchmark or reference rate for short-term U.S.-dollar interest rates. It is the rate of interest at which banks borrow funds from other banks in the London interbank market. Eurodollar futures are financial futures that have a final settlement value equal to the three-month LIBOR setting on the day of the contract's expiry.

periods and the associated high levels of volatility that the time-varying approaches produce significantly different results.[5]

The U.S. Federal Reserve has published a number of studies on extracting implied expectations from market interest rates. Gurkaynak, Sack, and Swanson (2002) examine a number of short maturity assets and determine that federal funds futures contracts dominate all other instruments in forecasting changes in the federal funds rate over horizons of several months, and that eurodollar futures perform better than other instruments for the longer horizons. Sack (2002) demonstrates how to extract policy expectations from these instruments using two models. The first model assumes stable risk premiums, whereas the second allows the risk premiums to be time-varying. The results show that, for horizons out to one year, the impact of time-varying risk premiums is limited and that "extracting policy expectations under the assumption of a constant risk premium may not be too misleading for shorter terms" (Sack 2002, 19).

The recent evidence in support of the EH at the short end of the yield curve is central to the model of interest rate expectations developed in this paper. When the EH is rejected, there are two possible reasons. The first is that longer-term interest rates have provided accurate measures of market expectations, but that the expectations have proved to be inaccurate (expectational errors). The second possibility is that the risk premiums assigned by the market to longer-term rates is not constant, but rather varies over time. The more recent studies tend to have a shorter horizon and focus on a more current time period (the 1990s to the present), encompassing a time of generally increasing central bank transparency. These two factors may have helped to reduce expectational errors, allowing the EH to hold. Recent changes by the Bank, including an increased level of transparency and the implementation of its FAD policy, may have similarly resulted in market participants being able to formulate more accurate expectations. The increased level of transparency helps to ensure that market participants are aware of the Bank's view of the economy, and the implementation of the FAD has removed a large amount of uncertainty as to the actual timing of changes in the policy rate. If these changes have helped to reduce expectational errors in the Canadian market, then an expectations-based model may now be accurate over shorter time horizons. This proposition will be tested by evaluating the EH using a number of different short-term instruments. If it can be demonstrated that the EH does indeed provide a reasonable characterization of the term structure of short-term interest rates, then a model based on this hypothesis should be able to produce accurate measures of near-term expectations.

5. Examples of these crisis periods include the 1992 downgrade of Canada's foreign currency debt, the 1994 peso crisis, and the 1995 Quebec referendum.

3. Definition and Selection of Instruments

It is possible to use a variety of marketable instruments to extract market expectations over the short term. The periods initially focused on in this paper are one and three months, so the instruments first examined, in terms of maturity, are all three months and under. Specifically, the choice set of instruments includes treasury bills, schedule "A" bankers' acceptances (BAs), term purchase and resale (repo) agreements, overnight interest rate swaps (OIS), the 30-day overnight futures contract (ONX), and foreign exchange forward implied rates. The current overnight target rate is also used as an input. Section 3.1 defines and describes these various instruments.

3.1 Definition of instruments

3.1.1 The overnight interest rate

The overnight interest rate market is the market in which funds can be borrowed or lent between market participants for a term of one business day. The rate at which these transactions are conducted is referred to as the overnight rate and is quoted on an actual/365-day count basis. Although this rate is determined by the demand and supply conditions in the market, it is tightly linked to the Bank's target rate.

The Bank conducts monetary policy by setting the target for the overnight rate. The target rate is the midpoint of an upper and lower limit for the overnight rate. The range between the two limits is referred to as the operating band and is currently set at 50 basis points. The upper limit is the Bank Rate, which is the rate that the Bank charges to extend an overdraft loan to financial institutions overnight. The lower limit of this operating band is the rate of interest that the Bank pays to participating institutions that have surplus settlement balances at the end of the day. This arrangement effectively discourages transactions in the overnight market at rates outside of this band.

The overnight funding rate fluctuates around the target rate. Since the Bank implemented the FAD policy in November 2000, the overnight rate has tracked the target rate extremely closely (Figure 1).

Figure 1: Realized vs Target Overnight Rate

As Figure 1 shows, the relationship between the overnight (ON) rate and the target rate has been very tight, but it has become even closer since the implementation of the FADs. This relationship was tested by comparing the sum of squared deviations (SSD) over the two periods.[6] In the pre-FAD period, the SSD was 0.09 basis points, whereas in the post-FAD period it fell to 0.01 basis points.

3.1.2 Treasury bills

Treasury bills are short-term obligations of the Government of Canada. They are issued regularly for terms of three months, six months, and one year,[7] issued at a discount, pay no coupon, and mature at par. Prices are calculated on a simple interest basis and an actual/365-day count basis is used. Although activity in the treasury bill market has declined in recent years as a result of reduced primary issuance, the market remains quite liquid, with an average daily trading volume of approximately $4.24 billion.[8]

---

6. Calculated using the formula $SSD = \frac{\sqrt{\sum (actual - target)^2}}{\sqrt{n}}$.

7. Shorter maturity bills are issued for cash-management purposes at irregular intervals.

8. The outstanding treasury bill stock fell from a high of $170.5 billion in April 1995 to a low of $78.7 billion in December 2000. It was $104.4 billion at year-end 2002. The volume for both treasury bills and BAs is based on 2002 Investment Dealers' Association statistics.

The relationship between treasury bill yields and the overnight financing rate should be very tight, because investors have a choice between investing for a fixed term in treasury bills or reinvesting their funds daily at the overnight rate. Several idiosyncratic factors exist in the treasury bill market, however, that introduce potential problems in their use as a predictor of future overnight rates. Specifically, in the very short end of the market, supply and demand imbalances frequently result in treasury bills that have under two months to maturity trading below the actual overnight target.[9] This apparent mispricing can be difficult to arbitrage away, as financing short positions in the repo market is problematic.[10]

3.1.3 Schedule I bankers' acceptances

Schedule I BAs are tradable short-term corporate obligations that are backed by a line of credit and are guaranteed by the accepting banks. Although they can be issued for any maturity, BAs are typically issued for terms of one, two, three, six, and twelve months, with the majority of issuance concentrated at three months and under. Because they represent a corporate credit, BAs typically trade at a positive yield spread to government debt (they have a positive credit spread). The short term to maturity and the fact that they are guaranteed by a Schedule I bank, however, cause the credit spread to be both small and stable over time. As with treasury bills, BAs are issued at a discount and are priced using an actual/365-day count basis. BAs have recently represented one of the most liquid and observable instruments in the money market, with an average daily trading volume of approximately $6.3 billion.[11] While there is no repo market for BAs, which makes them impossible to sell short, issuance is frequent enough and large enough that any significant overpricing tends to be arbitraged away. This has resulted in BAs having a historically high correlation with the average overnight funding rate.

3.1.4 Term general collateral purchase and resale agreements

Banks use the overnight general collateral (GC) market to borrow and lend money on a daily basis, using Government of Canada securities as collateral. The term "general collateral" refers to the fact that the collateral pledged does not have to be a specific bond, and therefore the GC market is not influenced by a particular issue that is in short supply. As a result, these transactions are typically done at levels very close to the overnight target rate.

---

9. This is an example of how preferred habitat can lead to a negative term premium.
10. Selling a treasury bill short involves selling a bill that you do not own, and then borrowing that bill back in the repo market to make delivery. For structural reasons, very few short maturity bills are available to be borrowed in the repo market. This discourages market participants from establishing short positions in short maturity bills that may appear overpriced.
11. Investment Dealers' Association industry volume statistics.

It is possible for banks to borrow and lend for a longer term in this market, effectively locking in a funding cost for a fixed term (ranging out to one year). Because banks have a choice of either funding themselves on a daily basis (and assuming the daily refinancing risk) or locking in funding for a specific term, the term rates should represent an average of expected future daily rates, plus a term premium. This would seem to make term GC an ideal candidate for use in measuring overnight interest rate expectations. This market is relatively illiquid, however, with transactions occurring infrequently and with larger observed bid-offer spreads than other short-term assets. As well, largely because of the infrequency of trades, there is no historical database of term GC yields upon which to conduct analysis.

3.1.5 Overnight index rate swaps

The OIS is a fixed-to-floating interest rate swap that ties the floating leg of the contract to a daily overnight reference rate.[12] When an OIS matures, the counterparties exchange the difference between the fixed rate and the average CORRA rate over the time period covered by the swap, settling the trade on a net basis. Standard terms of one, two, three, six, and twelve months are available, but swaps can also be tailored to specific maturities. The fixed quote on an OIS should represent the expected average of the overnight target rate over the term of the agreement. The OIS market is relatively new, however, and there is not yet a sufficiently long database of historical yields to permit robust empirical analysis.[13]

3.1.6 Overnight repo rate futures (ONX) contracts

The ONX contract is modelled after the federal funds futures contract in the United States, with each contract representing the expected weighted-average overnight rate for a specific month (as measured by CORRA). Since the contracts are priced based on the target overnight rate, they have no credit-risk component to their yields. As well, since they are futures contracts, short positions are relatively inexpensive to maintain.

These contracts represent a direct measure of what the CORRA rate is expected to average during a specific future month. Any risk premium should simply represent compensation for the level of uncertainty over future overnight rates. The federal funds futures contract in the United States has become the standard for measuring market expectations for future changes in the federal funds rate. A Federal Reserve study (Gurkaynak, Sack, and Swanson 2002) has shown that, since 1994,

12. The floating leg of an OIS swap is set to the CORRA (Canadian overnight repo rate average) rate, which is an overnight rate published by the Bank based on interdealer broker data.
13. Databases of historical OIS closes are kept by both Reuters and Bloomberg, dating from the middle of 2001.

these contracts have dominated all other market interest rates for predicting changes in monetary policy over horizons out several months. It would seem to be a reasonable assumption that a Canadian contract following the same specifications could also perform well. The contract is fairly new, however, and both open-interest and trading volumes are relatively small compared with other money market products. This makes the contracts difficult to rely on in isolation, although they are useful as a check for results obtained using other instruments.

3.1.7 Foreign exchange forward implied rates

It is possible to extract implied term interest rates from the foreign exchange (FX) forward market. Covered interest rate parity maintains that, for the no-arbitrage condition to hold in foreign exchange markets, the forward price of a currency must be related to the interest rate differential between the two currencies that are being quoted. Specifically:

$$F = X\left(\frac{1+r}{1+r_f}\right)^t, \tag{2}$$

where F is the current forward exchange rate, X is the current spot rate, r is the domestic interest rate for term t, r_f is the foreign interest rate for term t, and t is the time to delivery of the forward contract (in years).

In Canada, the spread between the spot and forward rates for term t is a function of the spread between the U.S.-dollar LIBOR and the Canadian-dollar equivalent interest rate over the same term, t.

FX forward rates are quoted in terms of forward points. These points represent the premium (or discount) of the forward rate to the spot rate. Given the forward exchange rate, the current spot rate, the foreign interest rate (LIBOR), and the time to delivery, it is possible to calculate the implied domestic interest, r. Forward-point markets out to terms of three months are very liquid and transparent, with an average weekly volume of approximately $36 billion (Bank of Canada 2002).

3.2 Selection of instruments

The instruments selected to test the EH must meet a number of criteria. They must be frequently traded, liquid instruments with a relatively large outstanding stock or open interest, have minimal idiosyncratic factors that affect their yield, and should be used by a variety of investors. A large variety of investors for a given security permits the capture of as broad a base of opinion as

possible. The prices of the selected instruments must also be readily observable, so that an accurate market quote can be consistently obtained. Finally, and equally importantly, a historical series of yields must be available to allow for robust empirical testing. Table 1 lists the various money market instruments and identifies their characteristics.

Table 1: Money Market Instruments

Instrument	Liquid	Large stock/ open interest	Observable quotes	Historical database
Treasury bills	Yes	Yes	Yes	Yes
Schedule I BAs	Yes	Yes	Yes	Yes
Term GC	No	Yes	No	No
OIS swaps	Yes	Yes	Yes	No
ONX contracts	No	No	No	No
FX forward	Yes	Yes	Yes	Yes

As the table shows, only treasury bills, BAs, and FX forward implied rates meet all of the criteria needed to test the EH. For this reason, the empirical testing will be restricted to those instruments.

As the graphs in Figure 2 show, the yields on the instruments selected closely track both each other and the realized average overnight rate over their maturity.

Figure 2: One- and Three-Month Market Rates

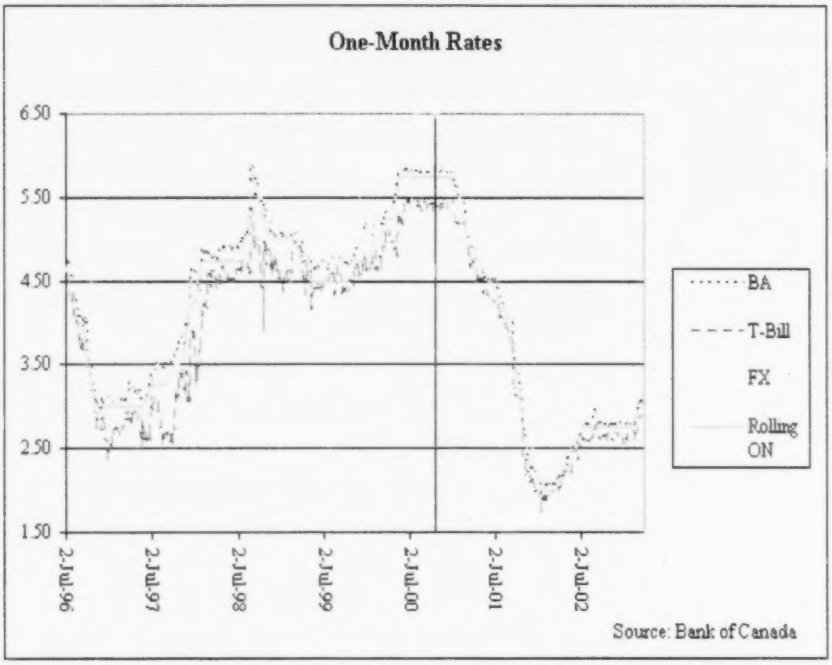

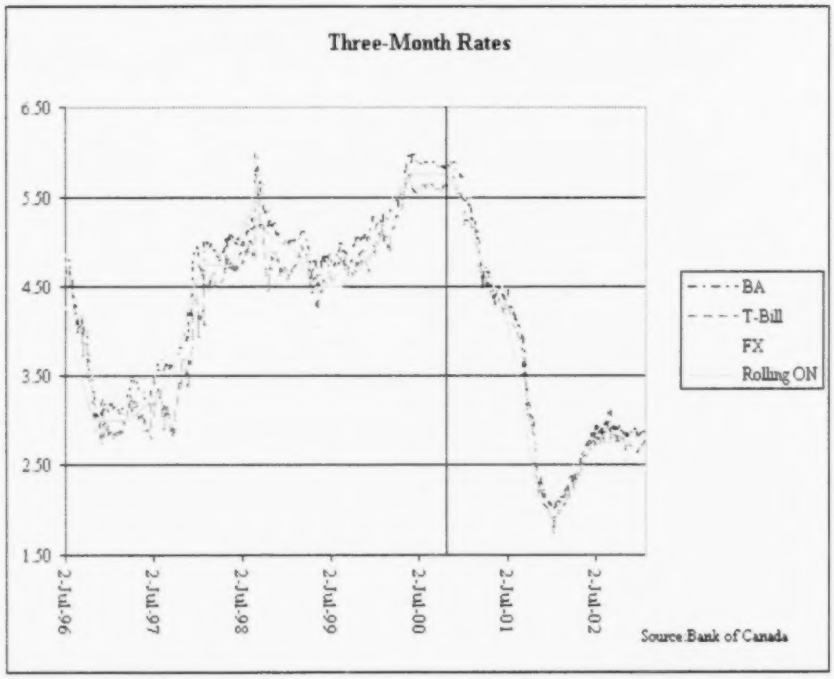

4. A Comparison of Pre- and Post-FAD Periods

The various instruments are examined over the pre- and post-FAD periods. For this study, the pre-FAD period covers from 2 July 1996 to 31 October 2000, and the post-FAD period from 1 November 2000 to 26 March 2003. Each sample is evaluated independently for each instrument. The results are examined to determine two points: whether one instrument dominates the others in terms of explanatory power, and whether the explanatory power of the instruments is changed by the implementation of the FAD regime.

4.1 Statistical summary of the periods

4.1.1 One-month assets

The graphs and statistics contained in Figure 3 show the ex-post excess returns earned for one-month BAs, treasury bills, and FX forward implied rates over the periods examined.[14] Two main points arise. First, while excess returns from BAs and FX implied yields appear to be slightly positive on average, treasury bill excess returns are larger in absolute magnitude and predominantly negative. Treasury bills, as direct obligations of the government, carry no credit risk. The positive excess for both BAs and FX implied yields is partially composed of a credit spread. The combination of no credit spread and a preferred habitat effect produces negative excess returns (a negative risk premium) for treasury bills. Second, all of the excess returns appear to have become less volatile in the post-FAD regime.

Figure 3: Excess Returns

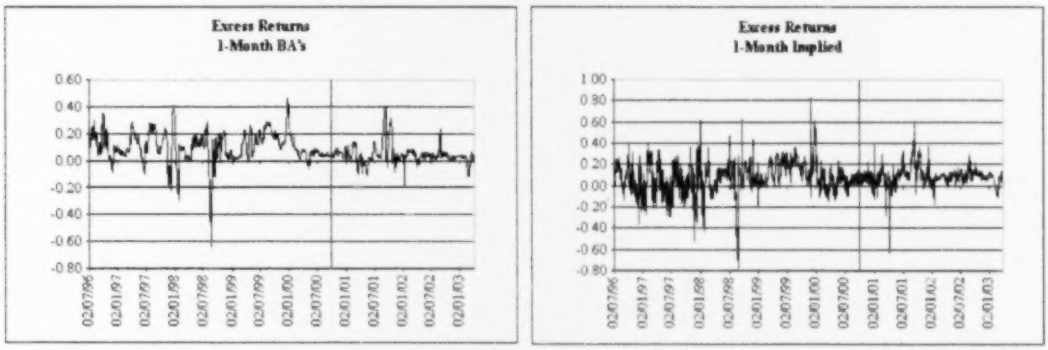

(continued)

14. The excess return is the difference in return earned by holding the term asset versus rolling the investment over daily at the overnight rate (the geometric average overnight rate over the period).

Figure 3 (continued): Excess Returns

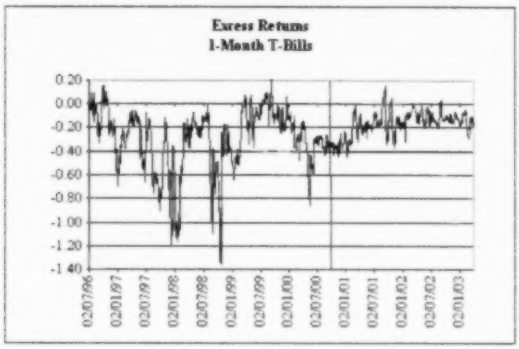

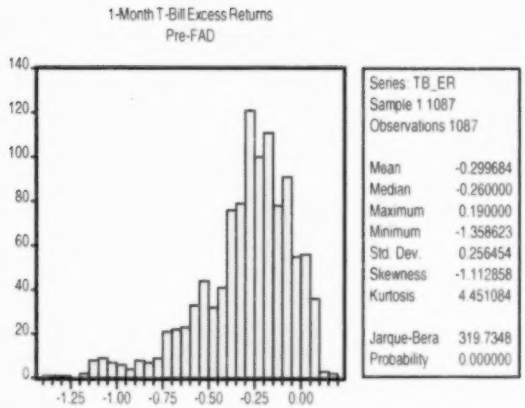

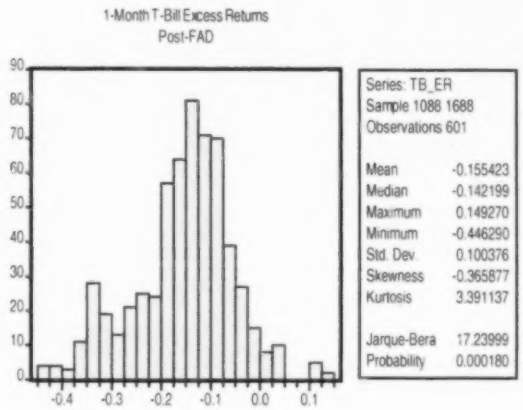

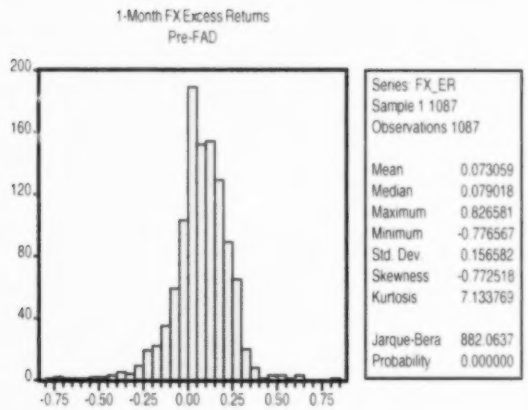

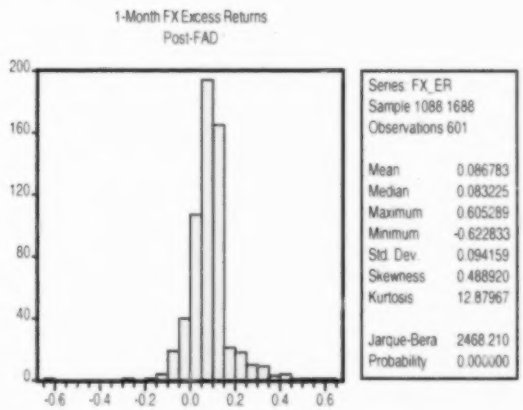

(continued)

Figure 3 (concluded): Excess Returns

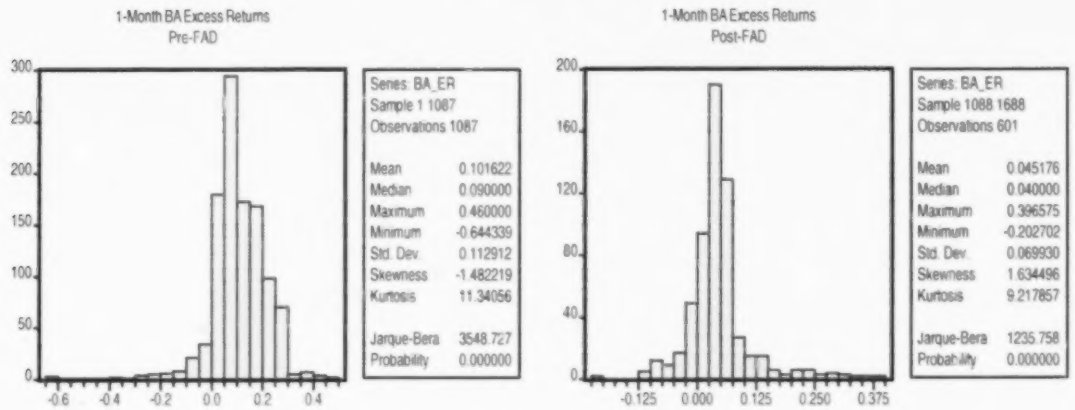

Table 2: One-Month Excess Return Correlations

Pre-FAD	BAs	T-bills	FX implied		Post-FAD	BAs	T-bills	FX implied
BAs	-	0.56	0.70		BAs	-	0.53	0.68
T-bills	0.56	-	0.44		T-bills	0.53	-	0.44
FX implied	0.70	0.44	-		FX implied	0.68	0.44	-

Figure 3 shows that, compared with treasury bills, BAs and FX implied rates have excess returns that are smaller in magnitude and less volatile in both the pre- and post-FAD periods. As well, the absolute values for the excess returns for BAs and treasury bills decrease in the post-FAD period. For BAs, the mean excess return falls from approximately 10 basis points to 4.5 basis points. For treasury bills, the mean moves from -30 basis points to -16 basis points. These differences are significant at the 99 per cent confidence level.[15] The results for the FX forward implied rates show the opposite behaviour, however, with the excess return increasing from an average of 7.3 basis points to 8.7 basis points. This difference, while relatively small, is still statistically significant.[16] Volatility of the excess returns (as measured by the standard deviation) falls significantly for all three assets in the post-FAD period. Table 2 shows that the correlation between the excess returns of the various assets does not change materially between the pre- and post-FAD periods. The correlations between treasury bills and the other instruments remain relatively low, highlighting the importance of security-specific factors in the pricing of one-month treasury bills.

15. The *p* values for a heteroscedastic *t*-test for both BAs and treasury bills are 0.00.
16. The *p* value for a heteroscedastic *t*-test for FX implied yields is 0.05.

Excess returns are a measure of the risk premium associated with the asset, so lower and less variable excess returns are indicative of smaller risk premiums (the negative excess returns for treasury bills highlight the riskless nature of that asset). The smaller, less variable excess returns in evidence in the post-FAD period are consistent with smaller risk premiums being required. Although we cannot firmly establish a causal link, these results do at least suggest that the implementation of the FAD regime may have increased the pricing efficiency and lowered the risk premium of the one-month assets examined.

4.1.2 Three-month assets

The graphs and statistics contained in Figure 4 show the same excess return information for the three-month assets.

Figure 4: Excess Returns

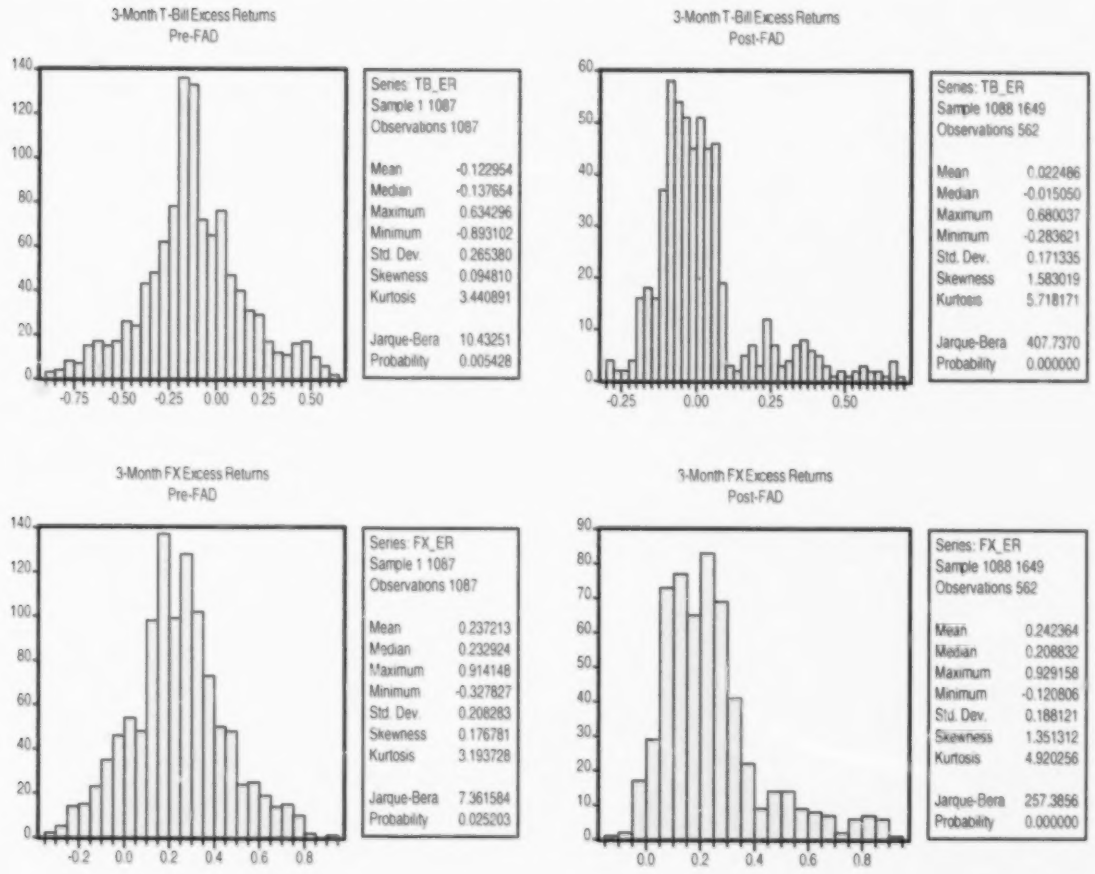

(continued)

Figure 4 (concluded): Excess Returns

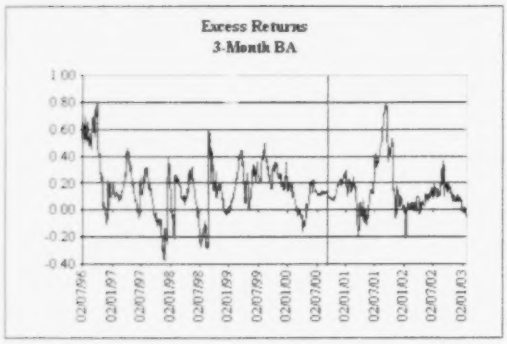

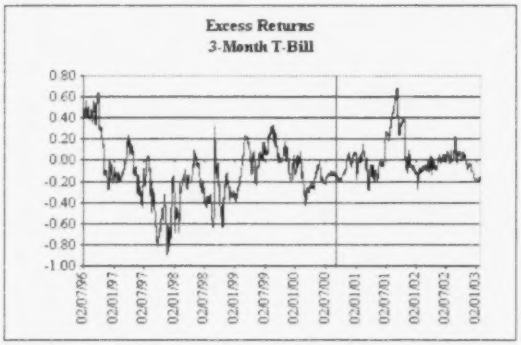

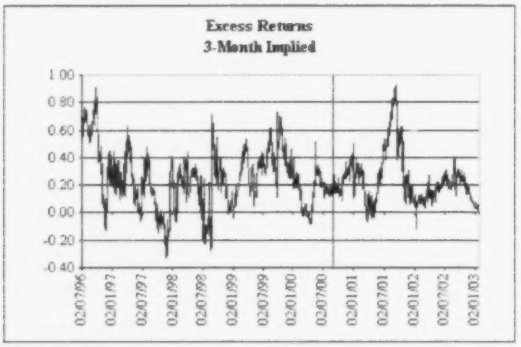

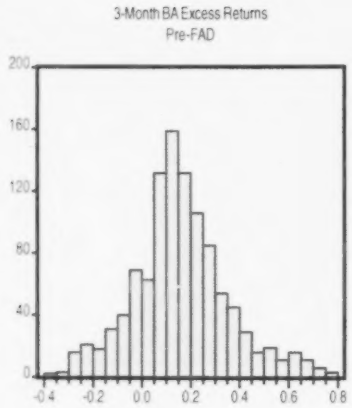

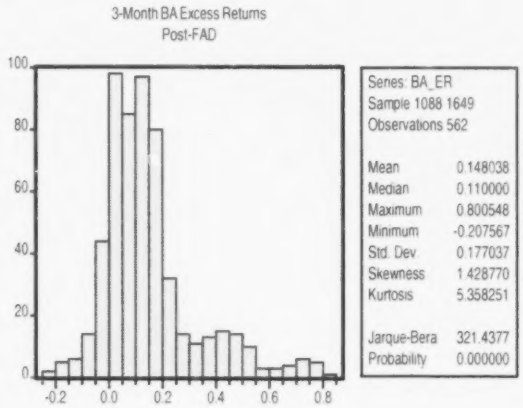

The summary statistics for the three-month assets show very different results than they did for the one-month assets. Whereas one-month assets generally show material decreases in both the absolute magnitude and variance of excess returns, for three-month assets only treasury bills show any statistically significant improvement. For BAs and FX forward implied rates, there is little apparent change in either the size or the volatility of the excess returns in the pre- and post-FAD

periods. Only treasury bills show a significant reduction in their risk premium after the implementation of the FAD regime.[17] The histograms also show that the distributions of excess returns look much less normal in the post-FAD period. The correlations of excess returns (Table 3), although high even in the pre-FAD period, become extremely high after the introduction of the FAD regime. The pricing behaviour of three-month assets appears to be far more homogeneous than it is for one-month assets, with security-specific factors playing a smaller role.

Table 3: Three-Month Excess Return Correlations

Pre-FAD	BAs	T-bills	FX implied		Post-FAD	BAs	T-bills	FX implied
BAs	-	0.88	0.93		BAs	-	0.97	0.97
T-bills	0.88	-	0.83		T-bills	0.97	-	0.95
FX implied	0.93	0.83	-		FX implied	0.97	0.95	-

The excess returns for the three-month assets probably behaved in this manner in the post-FAD period because of the terrorist attacks of 11 September 2001. The emergency easing that followed the attacks (i.e., outside of the normal FAD dates) would likely have had a large negative impact on the explanatory power of the various assets examined. The financial markets could not possibly have been able to anticipate such an event or the between-FAD emergency easing that followed, so excess returns would have been skewed artificially higher as a result of large expectational errors. The smaller sample size of the post-FAD period would exacerbate the problem. The effect would be larger for three-month assets, as their excess returns would be impacted for the three months prior to the emergency easing. This would result in 90 days of abnormal returns out of a sample of 562 days (16 per cent of the sample). In comparison, one-month assets would have only 30 days of abnormal returns (5 per cent of the sample). This conclusion is supported by the change in the skewness of the distributions. The histograms show that the distributions of excess returns become much more positively skewed in the post-FAD period. In section 4.2, it will be possible to adjust for the presence of this 11 September 2001 distortion when testing the explanatory power of the various assets.

17. The p values for a heteroscedastic t-test for BAs and FX implied yields are 0.27 and 0.62, respectively. For treasury bills, the p value is 0.00.

4.2 Testing the expectations hypothesis

The one-month statistical summary shows that, compared with treasury bills, BAs and FX implied yields have excess returns that are smaller in absolute magnitude and less variable in both the pre- and post-FAD periods. As well, the excess returns are generally smaller and less volatile after the implementation of the FAD regime. For three-month assets, treasury bills and BAs have the smallest absolute excess returns, and only treasury bills show a material decrease in the size of the excess returns in the post-FAD period. These summary statistics do not, however, indicate how effective these assets have been as predictors of the average realized overnight rate over their time to maturity, or whether this effectiveness has changed with the shift to the FAD regime. To do this, it is necessary to test the EH directly over the two periods, which means that it is necessary to define some forecasting equations to measure the explanatory power of the various assets.

According to the EH, the yield on a short-term asset from time t to time $t+r$, which will be denoted as $Y_{t,t+r}$, should be determined by the expected average overnight rate, ON_{t+j}, that will be realized over the period, plus a constant risk premium, α:

$$Y_{t,t+r} = E\left[\left[\prod_{j=t}^{t+r-1}(1+ON_j)-1\right]^{\frac{1}{r}}\middle|\Omega_t\right]+\alpha. \tag{3}$$

To estimate this relationship using an ordinary least-squares regression, equation (3) can be moved forward in time (using realized geometric average overnight rates) and rearranged to produce the following regression equation:

$$ON_{t,t+r} = -\alpha+\beta(Y_{t,t+r})+\varepsilon, \tag{4}$$

where $ON_{t,t+r} = \left[\prod_{j=t}^{t+r-1}(1+ON_j)-1\right]^{\frac{1}{r}}$ is the compounded return earned from rolling an investment over at the overnight rate.

According to the EH, β would be equal to one, α would be equal to the constant risk premium, and ε is the residual error. The coefficients that would be obtained by estimating equation (4) directly, however, would be dominated by the long-run relationships between the variables.[18]

18. To estimate equation (4) directly is essentially a test of cointegration between the asset's current yield (Y_t) and the subsequent realized overnight rates ($ON_{t,t+r}$). Preliminary tests of this model found strong evidence of cointegration, producing point estimates of β that ranged between 0.99 and 1.00 for all assets. This can be interpreted as at least weak evidence that the EH could hold over the long term.

Since the focus of this paper is on the short-run relationships (predicting the overnight rate over the next several months), the standard practice of stochastically detrending the equation by subtracting the current level of the overnight rate is followed. The regression equation becomes:

$$ON_{t,t+\tau} - ON_t = -\alpha + \beta(Y_{t,t+\tau} - ON_t) + \varepsilon. \quad (5)$$

This regression is estimated using historical yields for all three assets as the independent variable $Y_{t,t+\tau}$ for both sample periods. The estimation was performed using daily closing yields. This provided 1,087 observations in the pre-FAD period for both one- and three-month assets. In the post-FAD period, there were 601 observations for one-month assets and 562 observations for three-month assets.

To determine each asset's explanatory power in the post-FAD period, it was necessary to adjust for the impact of the emergency easing that occurred in response to the terrorist attacks of 11 September 2001. This was done by adding a dummy variable to equation (5) in the post-FAD regressions. The variable Θ_{sep11} is set to a value of zero if the excess return value for a given date is not impacted by the emergency easing, and to a value of one if the asset's excess return is impacted by the emergency easing (i.e., the easing occurred within t+30 days for one-month assets and t+90 days for three-month assets).

4.2.1 One-month assets

Table 4 reports the results of regressions run using one-month assets.

Table 4: One-Month EH Regression Results

Independent instrument	α	β	Θ_{sep11}	R^2	SEE	White test[a]	DW[b]	Chow test[c]	T-stat[d] $\beta = 1$	Term premium[e]
T-bill Pre-FAD	-0.02 (0.02)	0.04 (0.05)	n.a.	0.2%	0.14	1.2 (0.56)	0.12		19.7 (0.00)	
T-bill Post-FAD	-0.10 (0.01)	0.69 (0.07)	-0.22 (0.03)	67.9%	0.08	148.1 (0.00)	0.18	384 (0.00)	4.4 (0.00)	
BA Pre-FAD	0.08 (0.01)	0.78 (0.09)	n.a.	39.5%	0.11	59.1 (0.00)	0.12		2.4 (0.02)	
BA Post-FAD	0.04 (0.01)	0.99* (0.08)	-0.21 (0.05)	82.0%	0.06	125.1 (0.00)	0.23	50 (0.00)	0.1 (0.89)	3.9 bps (0.5 bps)
FX implied Pre-FAD	0.02 (0.01)	0.40 (0.06)	n.a.	18.1%	0.13	14.8 (0.00)	0.19		9.8 (0.00)	
FX implied Post-FAD	0.08 (0.01)	0.79 (0.08)	-0.22 (0.04)	66.4%	0.08	109.0 (0.00)	0.64	124 (0.00)	2.5 (0.01)	

a. The White test statistic indicates the presence of heteroscedasticity in the residuals. The *p*-value (in parentheses) represents the probability that the null hypothesis of no heteroscedasticity cannot be rejected.

b. The Durbin-Watson statistics are all well below 2, which indicates a serial correlation of residuals. The residuals are also heteroscedastic. The Newey-West HAC standard-error terms are shown in parentheses and are robust to both the serial correlation and heteroscedasticity of the residuals.

c. The Chow breakpoint test is used to determine whether there is a structural break in the estimated equations. In all cases, the *p* values are 0.00, which indicates a significant structural break.

d. The *p* values are shown in parentheses. Values of β that are not significantly different from one at the 10 per cent confidence level are indicated by an asterisk.

e. When the EH cannot be rejected, the term premium is obtained by re-estimating the equation and forcing the value of β to equal one.

The regressions show that BAs clearly emerge as the one-month instrument with the most explanatory power, having the highest adjusted R^2 value in both the pre- and post-FAD periods. Although the EH can be rejected in the pre-FAD period for all of the instruments tested (β is significantly different from one), in the post-FAD period it cannot be rejected when one-month BA yields are used. In samples where β is not significantly different from one, an estimate of the term premium can be made by re-estimating the regression equation and forcing the value of β to equal 1.0. In this re-estimation, the value of α that is produced for one-month BAs (3.9 basis points) can be interpreted as being the average risk premium over the post-FAD period. The standard error of the risk premium is very small (0.5 basis points), which indicates that the value is fairly stable over the entire sample. The dummy variable for the 11 September effect in the post-FAD period is negative and significantly different from zero in every case, which indicates the presence of abnormally positive returns during that period.

The EH can be rejected in both the pre- and post-FAD periods for treasury bills and FX implied rates. Treasury bills also have a surprisingly low explanatory power in the pre-FAD regime, with an adjusted R^2 of essentially zero. While this seems counterintuitive, an examination of Figure 3 (one-month excess returns) shows that, in the pre-FAD period, excess returns on one-month treasury bills were very large in magnitude and extremely volatile. It would appear that security-specific factors dominated the pricing of one-month treasury bills, leaving only a very tenuous relationship to the average level of overnight rates.

All regressions show high levels of both serial correlation and heteroscedasticity in the residual series. This is not surprising, given the specification of the regression equation. The regression attempts to estimate *ex-ante* expectations of future overnight rates using *ex-post* excess returns. By definition, the residual series is influenced by large expectational errors (when realized overnight rates are different from what was anticipated by the market). Figure 5 plots the residual series for one-month BAs, both for the full sample and for the post-FAD period. Residuals appear to have become smaller in the post-FAD period, with periods of general stability punctuated by spikes higher and lower. The spikes in the series marked with circles indicate episodes where the market was "surprised" by changes in the overnight rate. In these cases, the EH may have accurately measured expectations, but these expectations turned out to be incorrect. This behaviour leads to both heteroscedasticity (since the expectational errors lead to inconsistent variance) and serial correlation (the expectational errors affect excess returns for a number of consecutive periods). Nevertheless, the point estimates produced by the regressions are still unbiased estimates. The standard errors, however, have been calculated using a Newey-West covariance matrix and are robust to both serial correlation and heteroscedasticity.

Figure 5: Residual Plots

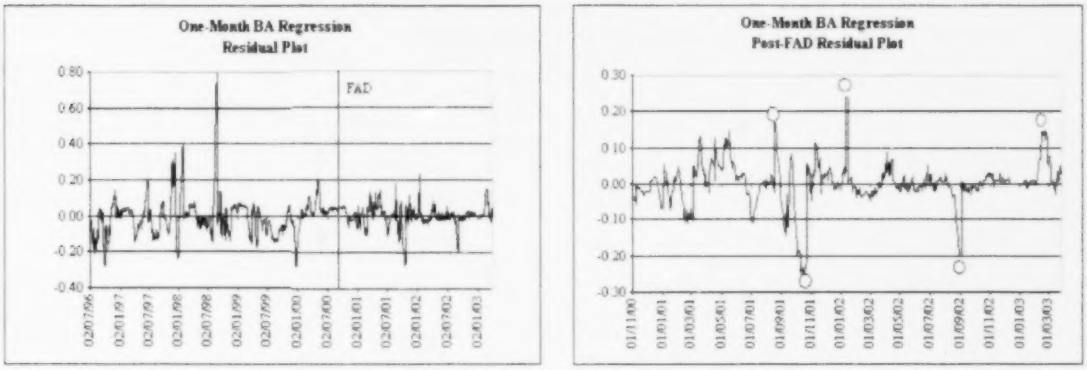

Generalized autoregressive conditional heteroscedasticity (GARCH) models are inappropriate for modelling the heteroscedasticity, because changes in the variance are not persistent. Since, by definition, the heteroscedasticity in the residuals is caused largely by expectational errors, prior levels of variance would be of little help in forecasting future levels (previous surprises don't provide information about future surprises).

Again, while no causality has been proven, the shift to the FAD regime appears to have been accompanied by a material increase in the pricing efficiency of one-month assets. The explanatory power of every asset increases substantially in the post-FAD period, and the EH could not be rejected in the post-FAD period when using one-month BAs. The Chow tests show that there is a definite structural break between the two periods. The proposition that the implementation of the FAD regime may be responsible for this improvement can be supported by observing the value of both the β and α coefficients over time. Figure 6 shows the evolution of the values of the estimates of both coefficients using a rolling window of 601 observations (the length of the post-FAD period). The sample is rolling, so the first observation after the vertical line covers one day post-FAD and 600 days pre-FAD. The process continues until the final observation, which is an entirely post-FAD sample. The graphs show that, after a brief period of time, as the rolling window moves beyond the implementation date, the value of β that the regression produced moves fairly quickly to converge towards the expected value of one. Once it reaches one, its value seems to become quite stable relative to prior periods. The values of α begin to move lower after the implementation of the FAD (apart from a spike higher around 11 September 2001), ending the period at the lowest level in the sample. The implementation of the FAD regime appears to coincide with both a move of the β estimates towards one and a move lower in the α estimates (risk premium). This provides additional support for the proposition that the increased evidence in favour of the EH in the post-FAD period is indeed the result of the new policy framework.

Figure 6: Rolling Coefficient Estimates

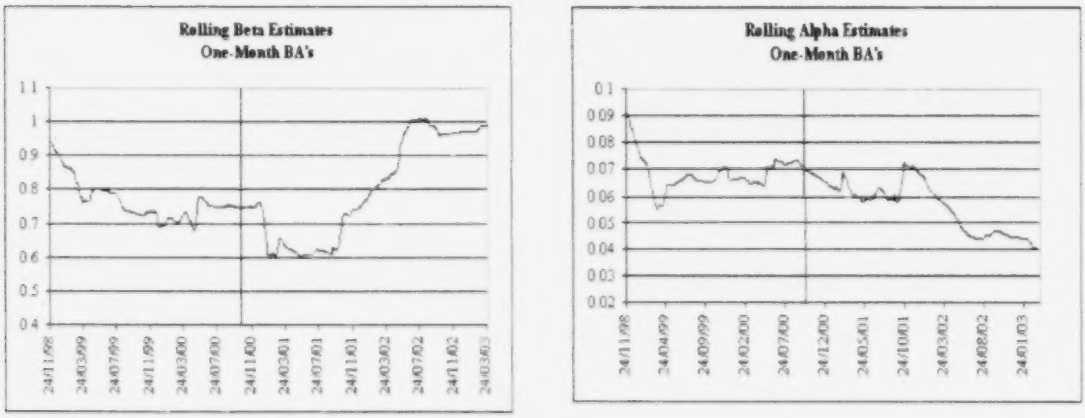

4.2.2 *Three-month assets*

Table 5 shows the results of regressions performed on the three-month assets.

Table 5: Three-Month EH Regression Results

Independent instrument	α	β	Θ_{sep11}	R^2	SEE	White test[a]	DW[b]	Chow test[c]	T-stat[d] $\beta = 1$	Term premium[e]
T-bill Pre-FAD	-0.09 (0.02)	0.59 (0.07)	n.a.	19.4%	0.25	16.9 (0.00)	0.02		5.8 (0.00)	
T-bill Post-FAD	-0.02 (0.01)	1.02* (0.06)	-0.37 (0.06)	84.9%	0.13	102.3 (0.00)	0.11	263 (0.00)	0.4 (0.71)	-1.9 bps (1.2 bps)
BA Pre-FAD	0.18 (0.02)	1.11* (0.08)	n.a.	51.8%	0.19	6.9 (0.03)	0.06		1.3 (0.19)	15.9 bps (1.5 bps)
BA Post-FAD	0.11 (0.01)	1.03* (0.07)	-0.38 (0.06)	83.4%	0.13	113.5 (0.00)	0.09	10.2 (0.01)	0.1 (0.96)	10.6 bps (1.3 bps)
FX implied Pre-FAD	0.23 (0.03)	0.95* (0.09)	n.a.	44.4%	0.21	13.6 (0.00)	0.14		0.5 (0.59)	23.7 bps (1.6 bps)
FX implied Post-FAD	0.20 (0.02)	1.02* (0.06)	-0.41 (0.07)	81.7%	0.14	73.1 (0.00)	0.16	8.7 (0.01)	0.4 (0.70)	19.7 bps (1.3 bps)

a. The White test statistic indicates the presence of heteroscedasticity in the residuals. The *p*-value (in parentheses) represents the probability that the null hypothesis of no heteroscedasticity cannot be rejected.

b. The Durbin-Watson statistics are all well below 2, which indicates a serial correlation of residuals. The residuals are also heteroscedastic. The Newey-West HAC standard-error terms are shown in parentheses and are robust to both the serial correlation and heteroscedasticity of the residuals.

c. The Chow breakpoint test is used to determine whether there is a structural break in the estimated equations. In all cases, the *p* values are 0.00, which indicates a significant structural break.

d. The *p* values are shown in parentheses. Values of β that are not significantly different from one at the 10 per cent confidence level are indicated by an asterisk.

e. When the EH cannot be rejected, the term premium is obtained by re-estimating the equation and forcing the value of β to equal one.

The three-month regressions yield results similar to those obtained using the one-month assets. The ranking of the assets in the pre-FAD period is the same as it was in the one-month regressions, with treasury bills performing the worst of all the assets and BAs performing the best. Again, all three assets show significant increases in their explanatory power, with all of the adjusted R^2 values over 80 per cent. Treasury bills move from last to first place, with an adjusted R^2 of almost 85 per cent, although the difference between first and last place is only 3.2 per cent. The 11 September dummy variable is negative and significantly different from zero in every case, which indicates the presence of abnormal excess returns around that event.

The EH cannot be rejected in the post-FAD period for any of the three assets. In every case, β is not significantly different from one. The EH also cannot be rejected in the pre-FAD period using either BA or FX implied rates. The three-month assets also appear to have had material increases in their explanatory power in the post-FAD period, and again the Chow tests show a structural break between the two periods for every asset. Although it is not possible to assign causality, the three-month assets also exhibit significant improvements in explanatory power and pricing efficiency in the post-FAD period.

The three-month regressions show similar levels of serial correlation and heteroscedasticity of residuals, as do the one-month assets. Figure 7 plots the residual series for three-month BAs in both the full period and the post-FAD period. As with one-month BAs, the post-FAD residual series appears to be generally more stable, although it is also marked by a number of spikes. Those spikes in the series that coincide with the surprises shown in the one-month series are again indicated with circles. As with the one-month assets, the residual series is influenced by relatively large expectational errors, leading to both heteroscedasticity and serial correlation. GARCH models are again inappropriate, but Newey-West standard errors are used for any hypothesis testing.

Figure 7: Residual Plots

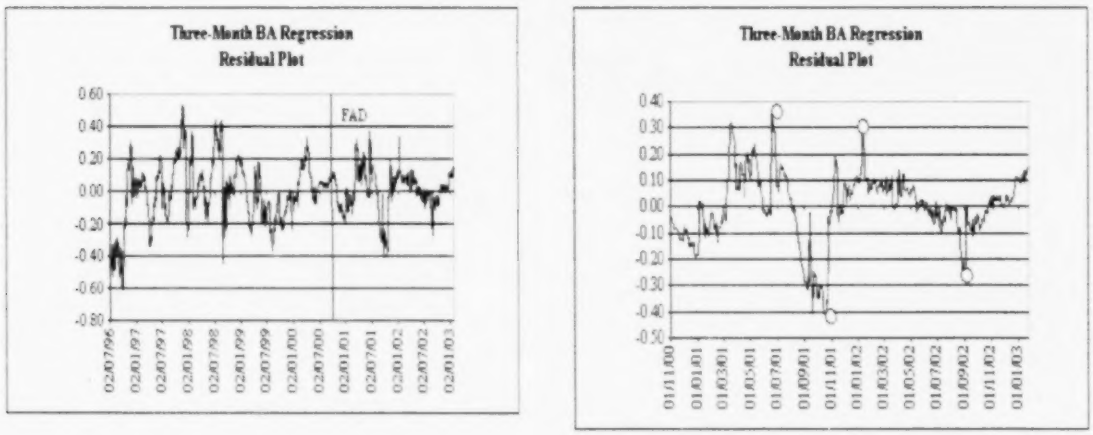

The same stability tests are performed to determine how the estimates of the β and α coefficients evolved over time for the three-month assets. Figure 8 shows a pattern similar, if somewhat less extreme, to that of Figure 5 for the one-month assets. After a brief adjustment period as the rolling window moves past the FAD implementation date, the value of β moves to converge towards a value of one, and the estimates of α start to trend lower, eventually settling at approximately 10 basis points, about 4 to 5 basis points lower than the pre-FAD level.

Figure 8: Rolling Coefficient Estimates

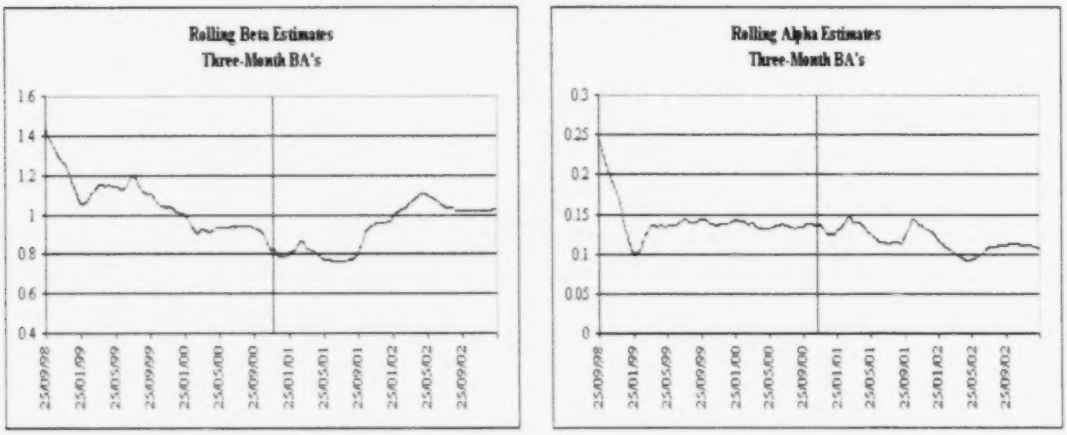

4.3 Final rankings and term premium estimates

Tables 6 and 7 rank the various assets in both the pre- and post-FAD samples. The assets are ranked in the order of their explanatory power (adjusted R^2 value). Also identified are the estimated values for their term premiums and whether the EH can be rejected.

Table 6: Pre-FAD Summary

Rank	One-month asset	R^2	Term premium[a]	Reject EH		Three-month asset	R^2	Term premium	Reject EH
1	BA	39.5%	10 bps	Yes		BA	51.8%	16 bps	No
2	FX implied	18.1%	7 bps	Yes		FX implied	44.4%	24 bps	No
3	Treasury bill	0.2%	-30 bps	Yes		Treasury bill	19.4%	-12 bps	Yes

a. If the EH cannot be rejected, the term premium is estimated by setting β equal to 1.0 and re-estimating the equation. If EH is rejected, then the average excess return is used.

Table 7: Post-FAD Summary

Rank	One-month asset	R^2	Term premium	Reject EH		Three-month asset	R^2	Term premium	Reject EH
1	BA	82.0%	4 bps	No		Treasury bill	84.9%	-2 bps	No
2	Treasury bill	67.9%	-16 bps	Yes		BA	83.4%	11 bps	No
3	FX implied	66.4%	9 bps	Yes		FX implied	81.7%	20 bps	No

In the one-month sector, BAs emerge as the asset of choice to measure implied expectations, having the highest explanatory power and the lowest term premium (in absolute values). In the three-month sector, treasury bills have the highest adjusted R^2 and the smallest term premium in the post-FAD period. The explanatory power of the three assets is very close, however, with all assets having an adjusted R^2 in excess of 80 per cent. As well, the EH could not be rejected for any of the three-month assets in the post-FAD period.

The inability to reject the EH runs counter to much of the published research on the topic; however, some key differences between this and previous work that has rejected the EH can account for this discrepancy. First, relatively little empirical work has been done on the Canadian market. The work that has been done generally looks at only treasury bills and examines a much earlier time period (Hejazi, Lai, and Yang 2000 examine the period from 1960 to 1995). One of the main propositions of the current paper is that the shift to a FAD regime, as well as generally increased transparency, may have helped to reduce expectational errors and improve the efficiency of the pricing of short-term assets, allowing the EH to hold. The second difference is that much of the work that rejects the EH in the U.S. market focuses on assets of a longer term-to-maturity, often testing long bond rates (e.g., Campbell 1995; Schiller, Campbell, and Schoenholtz 1983; and Fama and Bliss 1987). Recent work that focuses on the very short end of both the U.S. and European yield curve finds that the EH does hold, and that term rates are unbiased estimators of future average overnight rates (Longstaff 2000b; Gurkaynak, Sack, and Swanson 2002; Durre, Snorre, and Pilegaard 2003).

Strictly following the preceding results, we should use a yield curve built from a combination of one-month BAs and three-month treasury bills to extract implied forward rates. To keep consistency of instruments, though, and since we lose very little incremental explanatory power, BAs will be used for both the one- and three-month terms.[19]

The graphs in Figure 9 plot the actual versus the in-sample fitted values for the overnight rate. Both one- and three-month BAs are used as independent variables. As the graphs show, the fitted values for the overnight rate closely track the actual realized values. As one would expect, the relationship is tighter for the one-month horizon.

19. Consistency of instruments across maturities allows for easier interpolation between data points.

Figure 9: Actual vs Predicted Overnight Rates

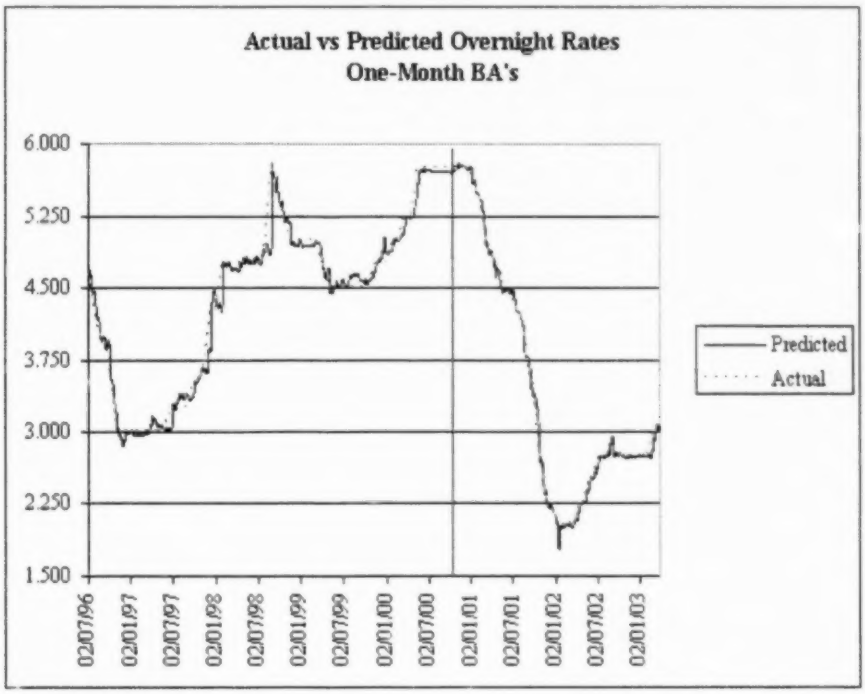

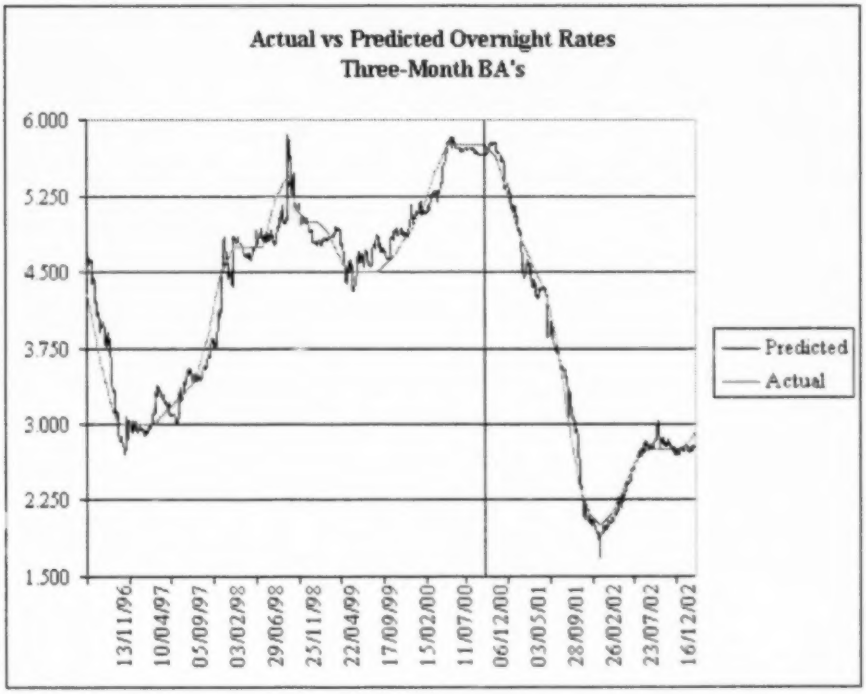

5. Expectations Over a Longer Horizon

5.1 Selection of instruments

The assets examined to this point have all been three months or less in maturity. This allows for the measurement of expectations out to three months in the future, covering at least the next two FADs. It is desirable, however, to have a way of measuring expectations for monetary policy beyond the next three months. Looking at expectations over a one-year horizon can be particularly helpful when trying to determine the expected timing of turning points in the interest rate path, or measuring the degree of cumulative tightening or easing expected up to a specific point in the future.

This measurement requires the use of assets that have a longer term-to-maturity than those examined thus far. The maturity of the assets selected needs to be at least as long as the time period over which expectations are to be measured. The assets selected need to meet the same criteria as those used at the shorter end of the maturity spectrum. They need to be frequently traded, liquid instruments with narrow bid/ask spreads and a large outstanding stock. Ideally, as with the shorter-term assets, the EH would not be able to be rejected.

Because BAs are the instruments used to construct the yield curve for maturities under three months, and because of the desirability to maintain consistency of instruments across the maturity spectrum, it would seem that longer-dated BAs would be an ideal candidate. There are problems with this approach, however. Although BAs are issued with terms-to-maturity of six and twelve months, these maturity tranches are quite illiquid.

There is an alternative to using longer-maturity BAs that still maintains the consistency of instruments: the 90-day bankers' acceptance futures contract (BAX), which trades on the Montreal Exchange. These contracts are similar to the eurodollar futures contracts in the United States. A BAX contract represents a notional amount of $1 million worth of three-month BAs, follows the International Monetary Market (IMM) dating standards, and converges to the three-month BA rate upon settlement.[20,21] The contracts currently represent one of the most liquid instruments in the Canadian money market, with average daily volume in the most active

20. The IMM is a division of the Chicago Mercantile Exchange. Financial futures contracts on this (and other) exchanges expire two business days prior to the third Wednesday of March, June, September, and December. These dates are referred to as IMM dates.

21. The BAX contract actually converges to the three-month BA rate, as measured by the daily survey of money market rates conducted by the Investment Dealers' Association of Canada.

contracts of approximately $1.2 billion.[22] While contracts exist that settle on the IMM dates for the next two years, volume and open interest drop off fairly sharply after the first three contracts.

The existence of these contracts allows the creation of "synthetic" BAs with terms-to-maturity of out to almost one year (using only the front three contracts). These synthetic BAs can then be used to measure expectations over a longer time horizon. To determine how effective these contracts are as indicators of expectations, however, it is necessary to test them using a version of the EH equation. Ideally, as with the shorter-term assets, these contracts will be shown to have a high degree of explanatory power and the EH would not be rejected. This test will also allow an estimation of what, if any, risk premiums exist in the various contracts.

5.2 Testing bankers' acceptance futures

This section will test whether BAX contracts represent unbiased predictors of future three-month BA rates. A version of equation (5) will be used, modified slightly to use the observed yields on the BAX contracts as the independent variable and the three-month BA rate at the contract's settlement as the dependent variable. A dummy variable will again be included to adjust for the impact of 11 September 2001. The equation will also contain a term to represent the time to settlement for the contract. This will allow the estimated risk premiums to change as the various contracts move closer to maturity. The specific equation is as follows:

$$3mBA_m - 3mBA_t = -\alpha + \beta_1(BAX_t - 3mBA_t) + \beta_2(d) + \beta_3(\Theta_{sep11}) + \varepsilon, \qquad (6)$$

where $3mBA_m$ is the three-month BA rate at the time of the front BAX contract's settlement, $3mBA_t$ is the current three-month BA rate at time t, BAX_t is the BAX contract yield at time t, α is the risk premium, and d is the time to the contract's settlement (in days).[23]

The equation is estimated using daily data for both the pre- and post-FAD periods to determine whether the predicted change in three-month BA yields (BAX_t-$3mBA_t$) is an unbiased predictor of the actual change ($3mBA_m$-$3mBA_t$). The null hypothesis is again $\beta = 1$. Table 8 shows the regression results.

22. The volume data are from the Montreal Exchange.
23. Days to settlement are divided by 100 in the estimation. This allows all the independent variables to be of similar magnitude.

Table 8: BAX Contract EH Regression Results

Independent instrument	α	β	Θ_{sep11}	d	R^2	SEE	White test[a]	DW[b]	Chow test[c]	T-stat[d] $\beta = 1$
Front BAX Pre-FAD	0.03 (0.03)	0.46 (0.16)	n.a.	0.11 (0.07)	15.9%	0.22	238 (0.00)	0.07		3.38 (0.00)
Front BAX Post-FAD	-0.05 (0.03)	1.53 (0.16)	-0.49 (0.06)	-0.24 (0.09)	66.5%	0.23	86 (0.00)	0.13	268 (0.00)	3.31 (0.00)
Front BAX Full Period	0.00 (0.02)	0.98* (0.15)	-0.54 (0.07)	-0.11 (0.07)	45.6%	0.25	336 (0.00)	0.09		0.13 (0.89)
Second BAX Pre-FAD	0.01 (0.06)	0.63 (0.09)	n.a.	0.08 (0.10)	35.8%	0.33	70 (0.00)	0.04		4.11 (0.00)
Second BAX Post-FAD	0.26 (0.12)	0.95* (0.10)	-1.11 (0.15)	0.01 (0.19)	58.3%	0.57	86 (0.00)	0.04	366 (0.00)	0.50 (0.62)
Second BAX Full Period	0.12 (0.07)	0.82* (0.10)	-1.28 (0.15)	0.04 (0.10)	61.0%	0.45	47 (0.00)	0.04		1.80 (0.07)
Third BAX Pre-FAD	0.22 (0.09)	0.96* (0.07)	n.a.	0.10 (0.12)	55.8%	0.45	31 (0.00)	0.04		0.57 (0.57)
Third BAX Post-FAD	0.28 (0.19)	0.65 (0.12)	-1.57 (0.17)	-0.15 (0.25)	71.7%	0.64	199 (0.00)	0.03	510 (0.00)	2.92 (0.00)
Third BAX Full period	0.22 (0.09)	0.87* (0.09)	-1.68 (0.12)	0.03 (0.13)	72.9%	0.54	51 (0.00)	0.03		1.44 (0.15)

a. The White test statistic indicates the presence of heteroscedasticity in the residuals. The *p*-value (in parentheses) represents the probability that the null hypothesis of no heteroscedasticity cannot be rejected.

b. The Durbin-Watson statistics are all well below 2, which indicates a serial correlation of residuals. The residuals are also heteroscedastic. The Newey-West HAC standard-error terms are shown in parentheses and are robust to both the serial correlation and heteroscedasticity of the residuals.

c. The Chow breakpoint test is used to determine whether there is a structural break in the estimated equations. In all cases, the *p* values are 0.00, which indicates a significant structural break.

d. The *p* values are shown in parentheses. Values of β that are not significantly different from one at the 10 per cent confidence level are indicated by an asterisk.

The results of the regressions show that, for the entire sample period, the EH cannot be rejected at the 90 per cent confidence level for the first and third contracts, and at the 95 per cent confidence level for the second contract. The value of β_2 is significant only for the front contract, indicating that, for it, the risk premium declines as the contract moves closer to its settlement date. The estimate of α for the front contract is zero, which shows that the entire risk premium is captured in the estimate of β_2. For the second and third contracts, the value of the β_2 coefficients is not significantly different from zero, while the estimates of α are all positive and significantly different from zero, which indicates a risk premium for these contracts that does not vary significantly over the term of the contract. This makes intuitive sense, because pricing of the front

contract is heavily influenced by the fact that its yield must converge to the three-month BA rate over a relatively short period of time (between 1 and 91 days). The closer to settlement the contract is, the more certainty there is regarding its final value. For the second and third contracts, their longer time to settlement (92 to 181 days for the second and 182 to 273 days for the third) make them less influenced by this convergence. The predictive power of these contracts is relatively strong, with adjusted R^2 values ranging from 47 per cent to 73 per cent.

The results for the subperiods are much less satisfactory. The estimates of β are very volatile, with the EH being rejected in both subperiods for the front contract, in the pre-FAD period for the second contract, and in the post-FAD period for the third contract. This behaviour is likely due to a small-sample-size problem and large expectational errors (given the longer term of these assets). The path of interest rates within the two subperiods helps to demonstrate this problem. As Figure 10 shows, the pre-FAD period (to the left of the vertical line) is marked by generally rising interest rates, with BAX yields climbing from approximately 3 per cent to 6 per cent. In the post-FAD period, on the other hand, rates fall rapidly from 6 per cent to 2 per cent. It is possible that market expectations consistently underestimate the degree of tightening in the pre-FAD period, and consistently underestimate the degree of easing in the post-FAD period.[24] This is indicative of a small-sample-size problem, and over an entire interest rate cycle the directional forecast errors net out.

Figure 10: BAX Yields

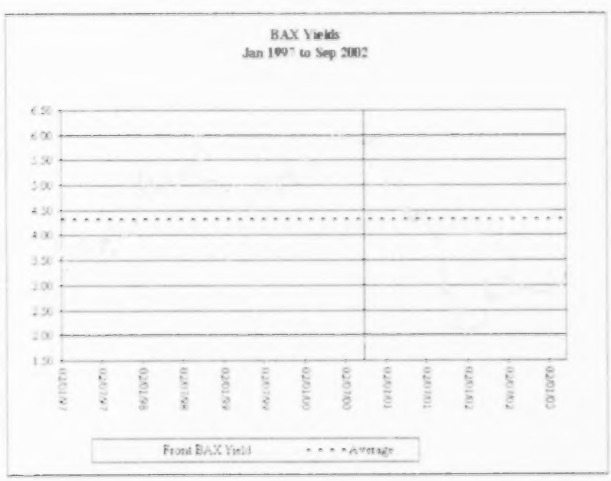

24. Market commentary over these periods is consistent with the proposition that the pace of changes in the overnight rate is in excess of that expected by market participants.

It is also possible that the shift to the FAD regime has less of an impact on the BAX market than it does on the shorter-maturity BA and treasury bill markets. The pricing of the one- and three-month instruments examined earlier is very sensitive to the specific timing of changes in the overnight rate. For a one-month asset, the exact day of a change in the overnight rate makes a significant impact on the pricing of the instrument. The BAX contracts, however, measure expectations over a longer horizon (from six to twelve months). This makes them less sensitive to the actual timing of the moves than to the general overall trend of monetary policy.

The full-period results of the regressions that are shown in Table 8 do support the hypothesis that, over an entire interest rate cycle, the front three BAX futures contracts are rational estimators of future three-month BA rates. The EH cannot be rejected at the 90 per cent confidence level for the first and third contracts, nor at the 95 per cent level for the second contract. These results are less robust than those obtained using one- and three-month assets, and it is possible that the maturity spectrum has moved far enough along that the time-varying component of the risk premium has become significant. Nevertheless, the results are sufficiently strong that they warrant the inclusion of the BAX contracts into the expectations model. The additional information content, combined with the fact that these contracts are widely used by market participants as a means of hedging and speculating on future changes in the overnight rate, outweighs the increased uncertainty generated by moving further along the time-to-maturity spectrum.

Given that the BAX contracts will be included in the expectations model, the regression equations are re-estimated to provide estimates of the values of the various risk premiums. Since the EH cannot be rejected, the equations are re-estimated with the value of β set to 1.0. As well, for the second and third contracts, the coefficient β_2 is set to zero. Table 9 shows the resulting term premium estimates.

Table 9: BAX Term Premium Estimates

Contract	Term premium estimate[a] (relative to three-month BA)	Total term premium[b]
Front	0.11 * (days to settle) (0.07)	11 bps + (0.11)(days to settle)
Second	15 bps (3 bps)	26 bps
Third	28 bps (4 bps)	39 bps

a. Standard-error estimates appear in parentheses.

b. The total term premium consists of the term premium estimated for each BAX contract relative to the three-month BA rate, plus the estimated three-month BA term premium (11 basis points).

The initial term premium values represent the premiums for the various BAX contracts relative to the value of a three-month BA at the time of the contract's settlement, whereas the total term premium values are relative to the overnight target. Figure 11 shows an estimated range of values (plus/minus one standard error) for the term premium on the various instruments.

Figure 11: Term Premium Ranges

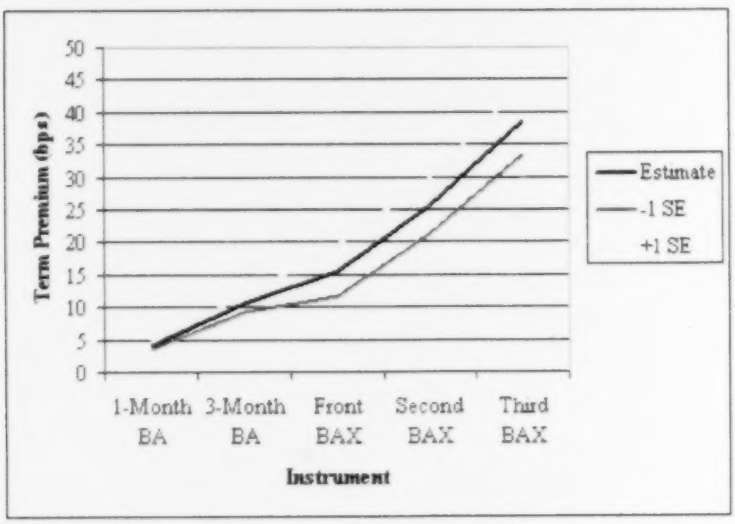

As Figure 11 shows, the estimates for the values of the term premiums range from very small and stable (one-month BAs, which have an estimate of 4 basis points and a standard error of 0.5 basis points) to fairly large and uncertain (the third BAX contract, which has an estimate of 39 basis points and a standard error of almost 6 basis points). Clearly, as the time horizon is extended, measures of expectations become increasingly uncertain. Estimates obtained using one- and three-month BAs appear to be quite precise. Once BAX contracts are introduced, these estimates become increasingly uncertain.

5.3 Generation of a spot BA yield curve

Once the observed yields on the various BAs and BAX contracts are adjusted for the presence of a term premium, the resulting yields can be used to build a spot BA yield curve with maturities extending out to approximately one year. This yield curve is generated by "rolling" together a series of BAs and BAX contracts to solve for the spot rate for any given maturity. The following example demonstrates the process, using price data from 17 December 2002:

One-month BA: 2.81% market yield less 4 bps term premium = 2.77%

Three-month BA: 2.87% market yield less 11 bps term premium = 2.76%

Front BAX (17 Mar 2003): 2.85% market yields less 21 bps term premium = 2.64%

Second BAX (16 Jun 2003): 3.03% market yields less 26 bps term premium = 2.77%

Third BAX (15 Sep 2003): 3.23% market yields less 39 bps term premium = 2.84%

A six-month BA can be replicated by purchasing a three-month BA and the front BAX contract, which settles on 17 March 2003. The payoff from this strategy is depicted by the following timeline:

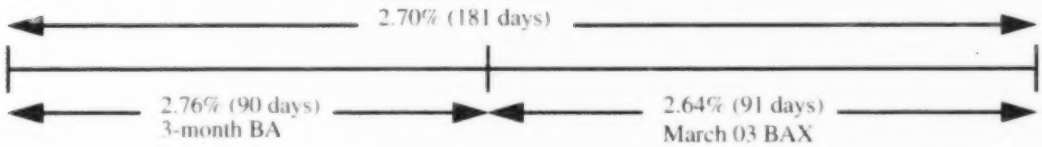

The six-month (181-day) rate is replicated by purchasing a three-month BA and rolling the investment over into another three-month BA at a guaranteed rate by using the BAX futures contracts. The effective rate is calculated as follows:

$$FV = (1.0276^{(90/365)})(1.0264^{(91/365)})$$

$$FV = 1.0133$$

$$181\text{-}day\ rate = FV^{(365/181)}$$

$$181\text{-}day\ rate = 2.70 \text{ per cent.}$$

This process can be extended to include the next two contracts, generating spot yields out to 15 December 2003.

6. Implied Forward Rates

The EH has been shown to provide a reasonably accurate description of the behaviour of yields in the Canadian money market. This implies that current market yields can be used to extract market expectations about the future level of short-term interest rates. Specifically, a spot yield curve can

be constructed using adjusted market rates for BAs and BAX contracts. This spot curve can then be used to calculate a series of implied forward overnight rates for various points in the future. These implied forward rates represent the market's expectation of the level of the overnight rate at a specific point in time.

Implied forward rates are essentially break-even rates; they represent what the level of a future interest rate would have to be to equate the holding-period returns of two different investment strategies. This is simply an extension of the EH, which states that the yield on a multi-period asset is the geometric average of expected future overnight rates, plus a risk premium.

The forward rate f at time a for period b can be expressed by the following equation:

$$f = \frac{(1 + z_{(a+b)})^{(a+b)}}{(1 + z_a)^a} - 1, \qquad (7)$$

where z is the spot interest rate for a given maturity.

Using this methodology, the spot yield curve can be used to calculate the implied forward short rate for any specific date in the horizon under examination. Since we have restricted the spot yield curve to assets with a time to maturity of approximately one year and less, it is only possible to calculate forward interest rates going out to twelve months. This period, however, covers the horizon that is of most interest to policy-makers. Expectations become far more uncertain extending beyond one year.

The following two examples show two ways of deriving market expectations using implied forward rates. The examples represent two different interest rate environments: one of expectations of rising overnight rates, and the other of expectations of stable rates.

The first step in this process is to create a spot yield curve from the observed yields on BAs and BAX futures contracts. Section 5.3 outlined the process by which the prices of BAX contracts can be used to determine longer-term BA rates. Spot yields for the one- and three-month BAs are combined with yields implied by the BAX contracts and an interpolation algorithm is used to construct a smooth spot yield curve.

The first example shows an environment in which expectations are for gradually increasing interest rates (Table 10).

Table 10: Market Yields, 8 May 2002

Instrument	Maturity	Yield (per cent)	Estimated term premium	Adjusted yield (per cent)
Overnight rate	9 May 2002	2.25	0	2.25
1-month BA	11 June 2002	2.34	4 bps	2.30
3-month BA	13 August 2002	2.55	11 bps	2.44
Front BAX	17 Jun to 16 Sep 2002	2.73	15 bps[a]	2.58
Second BAX	17 Sep to 16 Dec 2002	3.20	26 bps	2.94
Third BAX	17 Dec to 16 Mar 2003	3.78	39 bps	3.39

a. The front BAX contract has 40 days to expiry. The BAX term premium is calculated as 0.11 * 40 days. The result (4.4 basis points) is added to the three-month BA term premium of 11 basis points.

These adjusted yields are used to create a spot yield curve (Table 11).

Table 11: 8 May 2002, Spot Yield Curve

Term	Spot yield (per cent)
1-day - 9 May 2002	2.25
1-month - 8 June 2002	2.28
3-month - 8 August 2002	2.43
6-month - 8 November 2002	2.60
9-month - 8 February 2003	2.80

The final adjusted yield curve can be used to extract market expectations by using the implied forward overnight rates as of the upcoming FADs. In this example, the nine-month horizon spans six FADs. Table 12 shows implied overnight rates following each of the FADs and the associated probabilities for a change in the overnight rate.

Table 12: 8 May 2002, Implied Overnight Rates

Fixed announcement date	Implied overnight rate (per cent)	Probability of rate change
4 June 2002	2.50	100% of an increase to 2.50%
16 July 2002	2.54	16% of a further increase to 2.75%
4 September 2002	2.70	80% of a further increase to 2.75%
16 October 2002	2.86	44% of a further increase to 3.00%
3 December 2002	3.03	12% of a further increase to 3.25%
21 January 2003	3.35	40% of a further increase to 3.50%

As Table 12 shows, the market has fully priced in an increase in the overnight rate on the 4 June 2002 FAD from 2.25 per cent to 2.50 per cent. The 16 July 2002 FAD, however, has only a 16 per cent chance of a further hike from 2.50 per cent to 2.75 per cent priced in. The following FADs have progressively more tightening priced, with a total of 110 basis points of tightening priced in for the January 2003 FAD. The expected path of future overnight rates is depicted in Figure 12.

Figure 12: Implied Overnight Rates

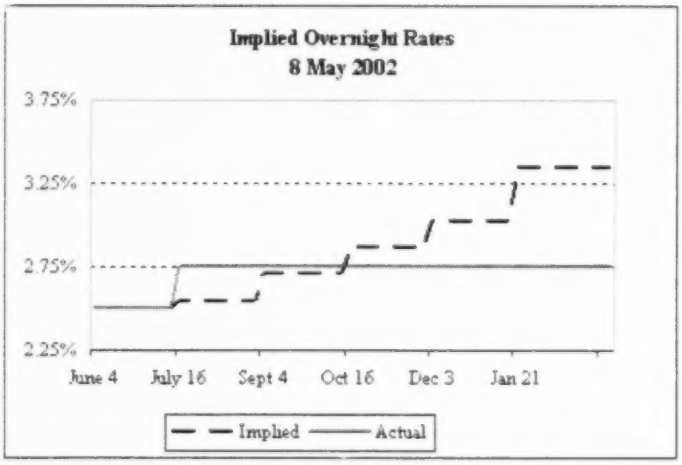

There are, however, some caveats to the probability calculations that appear in Table 12. The implied overnight rate gives the market's current expectation as to what level the overnight rate will be at on a given date. It does not provide information about the path that rates will take to

reach that level. The probability calculation assumes that the overnight rate can assume one of only two discrete values at the next FAD date. This is clearly an oversimplification, as there are some non-zero probabilities that the overnight rate could assume a wider range of possible values. Although the example above suggests that the market has priced in a 25 basis point increase in administered rates with 100 per cent certainty, it is also possible that expectations are split, with 50 per cent expecting no move and 50 per cent expecting a 50 basis point increase. Market convention, however, is to base probabilities on discrete intervals of 25 basis points, because, historically, the Bank has limited its changes to the overnight rate to increments of 25 (e.g., 25, 50, or 75 basis points). To transform the expected level of the overnight rate into a probability, it is necessary to assume that the current overnight rate (ON_t) can assume only two values after an upcoming FAD. Those values are ON_t with probability $1\text{-}P$ and ON_{t+1} with probability P. The value for P can then be solved as follows:

$$P = \frac{(f - ON_t)}{(ON_{t+1} - ON_t)}, \tag{8}$$

where f is the implied overnight rate at a point in the future.

The second example is from 17 December 2002 and represents a period of generally stable interest rate expectations (Table 13).

Table 13: Market Yields, 17 December 2002

Instrument	Maturity	Yield (per cent)	Estimated term premium	Adjusted yield (per cent)
Overnight rate	18 December 2002	2.75	0	2.75
1-month BA	20 January 2003	2.81	4 bps	2.77
3-month BA	17 March 2003	2.85	11 bps	2.74
Front BAX	17 Mar to 16 Jun 2003	2.85	21 bps[a]	2.64
Second BAX	17 Jun to 15 Sep 2003	3.03	26 bps	2.77
Third BAX	16 Sep to 15 Dec 2003	3.23	39 bps	2.84

a. The front BAX contract has 90 days to expiry. The BAX term premium is calculated as 0.11 * 90 days, which results in 9.9 basis points being added to the three-month BA term premium.

These yields are used to create a spot yield curve (Table 14).

Table 14: 17 December 2002, Spot Yield Curve

Term	Spot yield (per cent)
1-day - 18 December 2002	2.75
1-month - 18 January 2003	2.75
3-month - 18 March 2003	2.74
6-month - 18 June 2003	2.70
9-month - 18 September 2003	2.73

As Table 14 shows, once the various instruments are adjusted for term premiums, the yield curve is flat at approximately the current overnight rate of 2.75 per cent over the entire nine-month horizon. This is clearly indicative of market expectations for no change in policy. The actual implied forward rates are shown in Table 15.

Table 15: 17 December 2002, Implied Overnight Rates

Fixed announcement date	Implied overnight rate (per cent)	Probability of rate change
21 January 2003	2.75	0% chance of a change
4 March 2003	2.73	8% chance of a decrease to 2.50%
15 April 2003	2.65	40% chance of a decrease to 2.50%
3 June 2003	2.65	40% chance of a decrease to 2.50%
15 July 2003	2.73	8% chance of a decrease to 2.50%
3 September 2003	2.75	0% chance of a change

Figure 13: Implied Overnight Rates

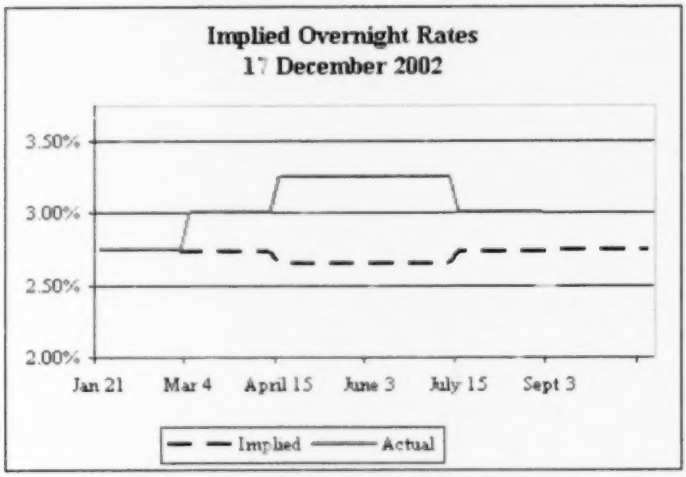

As Table 15 and Figure 13 show, market expectations on 17 December 2002 are for no change in the overnight rate at either of the next two FADs and relatively stable rates going forward.

While the forward-rate model produces specific results, some judgment must be applied in their interpretation. Table 15 shows that the implied overnight rate at the April and June FADs shows some probability of a decrease in administered rates. This result is inconsistent with other measures of expectations (such as survey date) and counter to the general market sentiment at the time (for stable rates). This decline in the forward rates over the April to June period is a result of the fact that the adjusted six-month spot rate (2.70 per cent) was lower than the adjusted three-month and nine-month rates (2.74 per cent and 2.73 per cent). This dip in the spot yield curve could be a result of the actual term premium for six-month yields (or the front BAX contract) being smaller than the estimated value that is used in the model.

The regressions performed in the previous sections showed that the term premium values have been relatively stable over time for the shorter-term assets. The one-month BA term premium was 4 basis points with a standard error of only 0.5 basis points, and the three-month BA term premium was 10.6 basis points with a standard error of 1.3 basis points. BAX contracts, however, had a relatively wider confidence band for their term premium estimates. For the front BAX contract, the term premium (in basis points) was estimated as 0.11 multiplied by the number of days to settlement, or approximately 10 basis points relative to the three-month BA in the 17 December example. The standard error of this coefficient, however, is 0.07. This gives a range of 3.6 to 16.2 basis points. As with any model, some judgment needs to be exercised, particularly when looking at expectations beyond the next couple of FADs. It is important to use other

instruments (such as the OIS market), as well as some of the more qualitative measures, as a check to support the results obtained from the implied-forward-rate model.

7. Conclusions

One of the primary methods currently used by the Bank to quantify market expectations for future changes in the overnight rate is to use an expectations-based model to derive implied forward rates. These implied forward rates are interpreted as the market's expectations for the future levels of short-term interest rates, and can provide both a point estimate and an implied probability of a future change.

The use of an expectations-based model assumes that the EH provides an accurate representation of the behaviour of money market yields in Canada. The results of this analysis show that, since FADs have been implemented at the Bank, the EH cannot be rejected for a number of short-term assets. BAs, in particular, show both a high degree of explanatory power and relatively small and stable risk premiums for all maturity tranches. As a result, one- and three-month BAs were selected as the primary input into the expectations model. The results for BAX futures contracts are not as robust as those for the BAs, reflecting their longer time-to-maturity and the increasing impact of expectational errors. Nonetheless, the results are strong enough to warrant their inclusion in the model.

The results of this analysis also show that the implementation of the FAD regime coincided with a significant improvement in the explanatory power of the various short-term assets examined. It was only in the post-FAD regime that the EH appeared to hold for the one- and three-month assets. Although no causality has been proven, it is possible that the generally higher level of transparency and reduced uncertainty surrounding the timing of changes in the overnight rate in the post-FAD period have reduced expectational errors in the pricing of money market assets. This inability to reject the EH for short time horizons is consistent with some of the more recent work conducted by both the U.S. Federal Reserve and the European Central Bank, and provides empirical justification for the use of an expectations-based model to derive market expectations.

Bibliography

Bank of Canada. *Bank of Canada Banking and Financial Statistics, January 2002*. Ottawa: Bank of Canada.

Campbell, J. 1995. "Some Lessons from the Yield Curve." *Journal of Economic Perspectives* 9(3): 129–52.

Cole, S. and W. Reichenstein. 1994. "Forecasting Interest Rates with Eurodollar Futures." *The Journal of Futures Markets* 14(1): 37–50.

Cox, J.C., J.E. Ingersoli, and S.A. Ross. 1981. "A Re-examination of Traditional Hypothesis about the Term Structure of Interest Rates." *Journal of Finance* 36(4): 769–99.

Durre, A., E. Snorre, and R. Pilegaard. 2003. "Estimating Risk Premia in Money Market Rates." European Central Bank Working Paper No. 221. Frankfurt am Main: European Central Bank.

Fama, E. and R. Bliss. 1987. "The Information in Long-Maturity Forward Rates." *American Economic Review* (77): 680–92.

Froot, K. 1989. "New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates." *Journal of Finance* (44): 283–305.

Gravelle, T. 1998. "Extracting Market Expectations at the Bank of Canada." A case study for the Bank of England's Centre for Central Banking Studies' *Workshop on Extracting Information from Financial Markets*, 15–19 June 1998.

Gravelle, T. and J. Morley. 1998. "Testing the Expectations Hypothesis: A Time-Varying Parameter Approach to Estimating the Term Premium." Bank of Canada. Unpublished manuscript.

Gravelle, T., P. Muller, and D. Stréliski. 1998. "Towards a New Measure of Interest Rate Expectations in Canada: Estimating a Time-Varying Term Premium." In *Information in Financial Asset Prices*, 179–216. Proceedings of a Conference Held by the Bank of Canada, May 1998. Ottawa: Bank of Canada.

Gurkaynak, R., B. Sack, and E. Swanson. 2002. "Market-Based Measures of Monetary Policy Expectations." Finance and Economics Discussion Series 2002–40. Washington: Board of Governors of the Federal Reserve System.

Hejazi, W., H. Lai, and X. Yang. 2000. "The Expectations Hypothesis, Term Premia, and the Canadian Term Structure of Interest Rates." *Canadian Journal of Economics* 33(1): 133–48.

Iimanen, A. 1995. *Understanding the Yield Curve*. New York: Salomon Brothers.

Longstaff, F.A. 2000a. "Arbitrage and the Expectations Hypothesis." *Journal of Finance* (55): 989–94.

———. 2000b. "The Term Structure of Very Short-Term Rates: New Evidence for the Expectations Hypothesis." *Journal of Financial Economics* (58): 397–415.

Paquette, J. and D. Stréliski. 1998. "The Use of Forward Rate Agreements in Canada." *Bank of Canada Review* (Spring 1998): 57–71.

Sack, B. 2002. "Extracting the Expected Path of Monetary Policy from Futures Rates." U.S. Federal Reserve Board Working Paper No. 2002–56.

Schiller, R. 1990. "The Term Structure of Interest Rates." In *Handbook of Monetary Economics*, Vol. 1, edited by B. Friedman and F. Hahn, 627–723. Amsterdam: North-Holland.

Schiller, R., J. Campbell, and K. Schoenholtz. 1983. "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates." *Brookings Papers on Economic Activity* 1: 173–217.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.***

2003

2003-25	Income Trusts—Understanding the Issues	M.R. King
2003-24	Forecasting and Analyzing World Commodity Prices	R. Lalonde, Z. Zhu, and F. Demers
2003-23	What Does the Risk-Appetite Index Measure?	M. Misina
2003-22	The Construction of Continuity-Adjusted Monetary Aggregate Components	J. Kottaras
2003-21	Dynamic Factor Analysis for Measuring Money	P.D. Gilbert and L. Pichette
2003-20	The U.S. Stock Market and Fundamentals: A Historical Decomposition	D. Dupuis and D. Tessier
2003-19	A Small Dynamic Hybrid Model for the Euro Area	R. Djoudad and C. Gauthier
2003-18	Technological Change and the Education Premium in Canada: Sectoral Evidence	J. Farès and T. Yuen
2003-17	Explaining and Forecasting Inflation in Emerging Markets: The Case of Mexico	J. Bailliu, D. Garcés, M. Kruger, and M. Messmacher
2003-16	Some Notes on Monetary Policy Rules with Uncertainty	G. Srour
2003-15	The Syndicated Loan Market: Developments in the North American Context	J. Armstrong
2003-14	An Index of Financial Stress for Canada	M. Illing and Y. Liu
2003-13	Un modèle << PAC >> d'analyse et de prévision des dépenses des ménages américains	M.-A. Gosselin and R. Lalonde
2003-12	The Macroeconomic Effects of Military Buildups in a New Neoclassical Synthesis Framework	A. Paquet, L. Phaneuf, and N. Rebei
2003-11	Collateral and Credit Supply	J. Atta-Mensah
2003-10	A Stochastic Simulation Framework for the Government of Canada's Debt Strategy	D.J. Bolder
2003-9	Bank Lending, Credit Shocks, and the Transmission of Canadian Monetary Policy	J. Atta-Mensah and A. Dib

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca